International Migration and Managerial Human Capital in Developing Economies

by

Emir Murathanoğlu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business and Economics)
in the University of Michigan
2024

Doctoral Committee:

Professor Dean Yang, Chair
Professor Achyuta Adhvaryu
Professor Anant Nyshadham
Professor Sebastian Sotelo

Emir Murathanoğlu
emirmur@umich.edu
ORCID iD: 0009-0009-2942-3780

ACKNOWLEDGEMENTS

Completing this dissertation would not have been possible without the guidance, feedback, and patience of my committee members Dean Yang, Ach Adhvaryu, Anant Nyshadham, and Sebastian Sotelo. My early conversations with Dean is what ignited my interest in a research agenda focused on international migration. His guidance on what is now the first chapter of my dissertation was invaluable for progressing the project from a crude idea to a complete research project. I am especially grateful for his encouragement at times when I lacked confidence in certain elements of the project myself. Sebastian, while the last member to join my committee, helped immensely in guiding the project to completion. His feedback has consistently helped me clarify my thoughts and recognize gaps in my arguments, providing me with ample opportunity to work through them. When I started working with Ach and Anant in my second year on what is now Chapter 3, I have never worked on an academic research project before. Working closely with them has taught much of what I learned about being an applied economist early on. I am indebted to all three for the immense effort they put in my mentorship and for providing an ideal to strive towards as an economist.

Along with my committee members, many other faculty members made it possible for me to arrive to this point. Chapter 1 of my thesis benefited greatly from discussions with John Bound, Charlie Brown, and Jagadeesh Sivadasan. I thank all of them for taking the time to attend my talks and provide feedback. Gaurav Khanna and Caroline Theoharides have been invaluable mentors and coauthors throughout my time in Michigan. It was a pleasure to learn from and work with them. Dominick Bartelme, as my third year paper reader, consistently encouraged me to think clearly and work on important questions. Thomas Buchmueller was instrumental in convincing me to apply for the joint business and economics program, a decision from which I never looked back. Beyond research, teaching alongside Tammy Feldman has set an example for me in what effective teaching looks like and helped develop my love for teaching. Finally, Faress Bhuiyan and Jenny Bourne deserves special mention for igniting my early interest in economics and ultimately putting me on the path to pursue graduate school.

My time in the program have been primarily shaped by the incredible people I met through the years who have been a constant source of support and friendship. I am especially grateful to Brick Rowberry and Iris Vrioni for the "random group" and the family dinners without which I don't see how I could have made it through the program; to Amy Ciardello for not storming out from these family dinners despite the endless econ-talk she had to endure for over half a decade; to Chris Hollrah and Emily Horton for always willing to host me at the end of my weekend walks and for teaching me how to drive; to Katherine Richards and Tyler Radler for having the (not so) secret Emir special in their wedding; to Nikhil Rao for furhering my love for many British comedy series and to Lea Bart for never judging me too hard when she witnessed me watch them for hours on end during the days of Covid. I would also like to thank Agostina Brinatti, Fernando Crepaldi, Çağla Ergül, Yaşar Ersan, Luis Espinoza, Andrea Foschi, Elird Haxhiu, Caitlin Hegarty, Thomas Helgerman, Billy Huang, Max Huppertz, Aaron Kaye, Will Mandelkorn, Nate and Faith (and Booker) Mather, Russell Morton, Zsigmond Palvölgyi, Hanna Onyshchenko, Paul Organ, Carolina Santos, Merve Sarıışık, Esteban Verdugo, and Triana Yentzen. Building these lifelong friendships is what I cherish the most from my 6 years in Michigan.

It is impossible for me to adequately express my gratitude for my family here, yet I will give it a try. First and foremost, mom and dad, none of this would have been possible without your lifelong and constant love and support. I love you both. Merve, my many conversations with you after setbacks in the program have been a constant source of comfort and encouragement. I am immensely lucky to have a sister like you. My soon-to-be brother Andrea, thank you for, along with seemingly making my sister very happy, organizing a trip with her to visit me when you realized that I was losing my mind during the dog days of the job market. I needed that. Grandma, I am sorry I have called you less than I should have over this busy last year. You asked me many times when I was going to be *finally* done with school. I am happy to report: "very soon!"

I dedicate this work to the memory of my grandfathers Cahit and Erol and my grandmother Berin. It is their enthusiasm and commitment towards the transformative power of education that has put me on the path I am on today. It is thanks to them I celebrate this milestone now.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

This dissertation contains three self-contained essays in the field of development economics that studies various aspects of international migration and managerial skill development in developing economies. Chapter I studies the effectiveness of temporary international labor migration as a shock-coping tool, using migration responses to typhoons from the Philippines as a case study. Chapter II studies the impacts of international migrant income on economic outcomes at migrant-origins. Chapter III studies the productivity impacts and within-firm allocation of a managerial soft skills training via a randomized controlled trial in a large ready-made garment firm in India.

Chapter I documents how international labor migration is used to cope with negative shocks, highlighting the role of search frictions and international macroeconomic conditions in shaping the migration response. Using administrative data on the universe of temporary migrant contracts from the Philippines, I show that typhoons increase international migration from affected municipalities. However, overseas wages of new migrant cohorts fall. The wage drop is driven by migrants leaving for low-wage countries and occupations, despite typhoon driven migrants being positively selected in terms of education. These patterns are consistent with congestion and search frictions in overseas contract markets: typhoons lower reservation wages and incentivize migrants to leave for lower paying overseas jobs. Strong international migrant demand dampens this response: typhoons lead to a larger migration increase without a proportionally large wage drop. As a result, the total migrant earnings response doubles when moving from median to 75th percentile migrant demand conditions. Similarly, households in typhoon affected regions receive more remittances when international migrant demand is high. These results suggest policies that increase the availability of overseas jobs in the wake of disasters can lead to substantial shock-coping gains.

Chapter II asks how does income from international migrant labor affect the long-run development of migrant-origin areas? We leverage the 1997 Asian Financial Crisis to identify exogenous and persistent changes in international migrant income across regions of the Philippines, derived from spatial variation in exposure to exchange rate shocks. The initial shock to migrant income is magnified in the long run, leading to substantial increases in income in the *domestic* economy in migrant-origin areas; increases in population education;

better-educated migrants; and increased migration in high-skilled jobs. 77% of long-run income gains are actually from domestic (rather than international migrant) income. We empirically demonstrate that these findings are not confounded by potential trade impacts of the same exchange rate shocks. A simple model yields insights on mechanisms and magnitudes, in particular, that 23.2% of long-run income gains are due to increased educational investments in origin areas. Improved income prospects from international labor migration not only benefit migrants themselves, but also foster long-run economic development in migrant-origin areas.

Chapter III studies the allocation and productivity consequences of managerial training via a randomized controlled trial among production line supervisors in a large ready-made garment firm. We designed a program using practices identified as productive in Adhvaryu et al. (2022d), and asked middle managers – who are directly above production line supervisors in the hierarchy – to recommend which of the supervisors they manage should be prioritized for training. We then randomized access to the program within these recommendation rankings. Productivity on lines managed by treated supervisors increased by 6-7% relative to control, but these gains exhibit substantial heterogeneity across middle manager ranking categories. Highly recommended supervisors experienced no productivity gains; the average treatment effect of training is driven entirely by low-recommendation supervisors. This was not due to a lack of information about baseline skills or about who would gain the most, nor to discrimination or favoritism along observable dimensions. Instead, consistent with the fact that supervisor turnover has large personal costs for middle managers in terms of labor substitution and onboarding, middle managers prioritized the retention impacts of training. Treated supervisors were 14% less likely to quit than controls over the study period, and this gain was driven by highly recommended supervisors. Heterogeneous returns and the unproductive allocation of costly training can thus help explain underinvestment in attenuating persistent within-firm gaps in managerial quality.

CHAPTER 1

When the Weather Turns: Coping With Shocks Through International Migration in the Presence of Search Frictions

1.0 Abstract

I document how international labor migration is used to cope with negative shocks, highlighting the role of search frictions and international macroeconomic conditions in shaping the migration response. Using administrative data on the universe of temporary migrant contracts from the Philippines, I show that typhoons increase international migration from affected municipalities. However, overseas wages of new migrant cohorts fall. The wage drop is driven by migrants leaving for low-wage countries and occupations, despite typhoon driven migrants being positively selected in terms of education. These patterns are consistent with congestion and search frictions in overseas contract markets: typhoons lower reservation wages and incentivize migrants to leave for lower paying overseas jobs. Strong international migrant demand dampens this response: typhoons lead to a larger migration increase without a proportionally large wage drop. As a result, the total migrant earnings response doubles when moving from median to 75th percentile migrant demand conditions. Similarly, households in typhoon affected regions receive more remittances when international migrant demand is high. These results suggest policies that increase the availability of overseas jobs in the wake of disasters can lead to substantial shock-coping gains.

1.1 Introduction

Every year, millions of migrants leave developing countries for temporary overseas employment. Access to such "guest worker" markets can potentially improve the ability of origin communities to cope with negative shocks by allowing a migration response, especially if shocks are spatially correlated and local risk-sharing networks fail.[1] Yet, evidence on whether and how such migration responds to negative shocks is limited. Given substantial policy interest in regulating and facilitating temporary labor migration from developing countries, it is important to assess its potential benefits as a coping strategy (United Nations, 2019).

In this paper, I study the impacts of typhoons on temporary international labor migration from the Philippines, one of the largest source countries of labor migrants. Theoretically, the impacts of typhoons on migration is ambiguous: natural disasters increase returns to migration by disrupting local economies; but they can also raise barriers to migration and worsen liquidity constraints, inhibiting potential migrants from bearing the up front costs of migration (Bazzi, 2017; Yang, 2008d). Further, temporary labor migration requires securing an overseas contract before leaving home. Therefore, the migration response depends on the ability to find overseas jobs *prior to* migrating. Two features of this market can hinder this. First, at prevailing overseas wages, the supply of willing migrants can exceed available overseas contracts (McKenzie et al., 2014; Mobarak et al., 2023a). Second, contracts are secured through decentralized search between migrants and many private recruitment agencies, where job-finding can be subject to search frictions. In this context, natural disasters can increase search for overseas contracts, but whether increased search induces additional migration depends on the availability of contracts and search frictions in job-finding. Beyond migrant flows, disasters can further impact migration outcomes by reducing the option value of continued search, incentivizing migrants to accept lower paying overseas contracts. Finally, the migration response can further depend on international macroeconomic conditions which influence international migrant demand and therefore the availability of overseas contracts.

To make progress in the face of this theoretical ambiguity, I estimate the impact of typhoons on international migration from 1597 Philippine municipalities over a decade. Beyond documenting the migration response, I aim to characterize how congestion and search frictions influence the migration outcomes, and how international migrant demand conditions mediate the migration response. First, I document that typhoons lead to increased out-migration, corroborating that migration is used as a coping mechanism. However, the migration response occurs through a larger share of migrants leaving for low-paying occupations and destination countries. Thus, new migrant cohorts have lower average wages following typhoons. Next, I show that the migration response is heterogeneous in migrant demand conditions facing a municipality: typhoons lead to more migration without a proportionally large drop in average migrant wages during periods of high international migrant

[1] For a review of the literature on the broader impacts of temporary migration on origin countries, see Bossavie and Özden (2023).

demand. I interpret these results through a model of migration with search frictions, which clarifies that typhoons can lead to increased search, but also lower the reservation wages of potential migrants; and that better migrant demand conditions can make the availability of overseas vacancies more responsive to the supply of potential migrants, leading to a stronger migration response.

To theoretically assess the potential effects of origin shocks, I begin by providing a simple model of migration choice with search frictions. A representative recruitment agency posts overseas vacancies and potential migrants search for overseas contracts to migrate. Due to search frictions, potential migrants may fail to match with a contract. If they match, they decide whether to accept the contract or keep searching based on the contract wage. In this environment, I consider a "typhoon" shock that decreases origin utility while raising barriers to migration. The shock spurs additional migration if the increased returns to migration dominates the higher barriers to migration. If the former dominates, the shock spurs migration by both increasing overseas job search and lowering reservation wages. This drop in reservation wages reduces average migrant cohort wages, even in the absence of selection or equilibrium wage responses abroad. Better international migrant demand conditions lower the vacancy posting costs of the recruitment agency. This leads to a larger migration response to shocks as recruitment agencies are more responsive to changes in the supply of willing migrants due to origin shocks. I assess these predictions empirically.

To measure migration outcomes, I use a comprehensive administrative dataset of new migrant contracts. The dataset covers the universe of new migrant contracts from the Philippines, with about 4 million migrant contracts between the analysis period of 2007 to 2016. It includes information on contract wage, duration, destination country, and occupation along with migrant demographics (sex and age) and municipality of origin. This detailed information allows for a comprehensive account of migration responses to disasters at a geographically granular scale. Beyond migrant flows, I can capture important margins of adjustment such as changes in contract wages or share of migrants leaving for different occupations and destinations. This level of detail on migrant origins, destinations, and contracted wages is generally not available with more common data sources such as censuses or surveys.

To measure typhoon exposure, I construct a continuous measure of annual exposure using meteorological data. The measure factors in the number, intensity, and impact-area population of typhoons hitting a municipality in a given year. The meteorological nature of measurement ensures that the measure is not prone to error or bias due to misreporting and that the measurement standards are consistent across space and time. I validate the measure by demonstrating that it predicts government estimates of typhoon damages and casualties, along with drops in nightlight intensity. I exploit within-municipality variation in typhoon exposure to estimate the impacts of typhoons on migration.

Armed with the migration and typhoon measures, I first document that typhoons cause an immediate and prolonged increase in temporary labor migration. A one standard deviation typhoon exposure increases the annual out-migration rate from Philippine municipalities by 1.2 and 1.4 per 10,000 in the short-run (1-2 years after exposure) and medium-run (2-3 years after exposure) respectively, corresponding to 3.7% and 4.3% of the mean migration rate. These results suggest that migration can be used as a coping mechanism and that excess supply of migrants or worsening liquidity constraints do not impede an average migration response.

However, the increase in migration is accompanied by a larger share of migrants leaving for lower paying destination countries and occupations. A one standard deviation typhoon exposure increases the share of migrants leaving for the lowest paying countries by 1.4 percentage points (1.9% of mean) and the share of migrants leaving for the lowest paying occupations by 0.8 percentage points (1.2% of mean) in the short-run. This *downgrading* is not driven by changes in migrant cohort demographics. In fact, using complementary survey and census data, I find that the educational attainment of migrant cohorts increase after typhoons, suggesting typhoon-induced migrants are positively selected in terms of education. Due to occupation and destination downgrading, average and median migrant cohort wages fall by 0.27% and 0.43% for each percent increase in migration in the short-run. The drop in new migrant wages leads to substantially lower *total migrant earning* response than what would take place without the wage drop ($28.5 million lower, corresponding to 22% of the average migrant earning response). The contemporaneous fall in migrant cohort wages and increase in migrant educational attainment is consistent with search frictions causing migrants to lower their reservation wages and possibly to direct search towards foreign markets with lower pay, but a higher likelihood of securing a job.

Next, I provide evidence that the ability to migrate in response to shocks is mediated by international migrant labor demand conditions facing a municipality.[2] I construct a plausibly exogenous municipality-year level shift-share proxy of international migrant demand by combining destination demand shifters (GDP shocks) with spatial variation in Philippine municipalities' exposure to destinations using baseline (1992-1997) migration shares. The identifying assumption is that the year-to-year variation in destination country GDP per capita is as-good-as-random from the perspective of individual Philippine municipality mi-

[2]While previous work has studied the direct effects of destination "pull factors" such as GDP and other proxies for labor demand, how demand conditions mediate responses to origin shocks is new to the literature to the best of my knowledge.

gration decisions. Using this proxy, I assess whether the migration response to typhoons is heterogeneous by contemporaneous demand conditions facing a municipality. I undertake inference using a randomization inference procedure following Borusyak and Hull (2020).

I find that the migration response to typhoons are substantially larger during periods of high migrant demand. The migration response to a one standard deviation typhoon exposure is essentially null during 20th percentile migrant demand conditions, but 6.7% during 80th percentile demand conditions. Further, the larger migration response during periods of high migrant demand is not accompanied by a correspondingly larger drop in migrant wages. Combining the larger migration response and smaller relative migrant wage drop, the total migrant earning response to typhoons is 2.3 times larger in 80th percentile migrant demand conditions compared to the median. The more muted wage response is partially driven by lower occupational downgrading during better demand conditions. Overall, these findings suggest that conditions and policies that increase the availability of overseas contracts in the wake of disasters can lead to significant shock-coping gains for affected regions. Further, the substantial heterogeneity by contemporaneous macroeconomic conditions suggests caution when generalizing from responses to individual shocks. Disasters that correspond to a period of extreme demand conditions may lack external validity (Rosenzweig and Udry, 2019).

Finally, given international migration aids origin shock-coping primarily through remittances sent home, I study remittances using household survey data from the Family Income and Expenditure Surveys. I find that typhoons cause a 6.7% and 4.5% increase in per capita remittances in the short-run and medium-run, respectively. Consistent with the migration results, remittance response to typhoons are larger during periods of high migrant demand, with a one standard deviation migrant demand improvement leading to a 67% increase in the remittance response to typhoons at the mean. To assess the importance of *new migration* in the remittance response, I combine the migration rate, migrant wage, and remittance estimates. 15%-19% of the remittance increases following typhoons can be plausibly attributed to new international labor migration, suggesting a substantial role for new migration in the aggregate remittance response.

The Philippines provide an ideal context for this study for several reasons. International labor migration is exceedingly common and many other prominent migrant origin countries have adapted migration programs with similar features to the Philippines, partially based on its perceived success (Theoharides, 2018a; Mishra and Rajan, 2010; Asis and Aguinas, 2012). Further, Philippine migrants leave for a wide variety of occupations and destinations which leads to significant migrant wage dispersion. This allows migrants to act on their reservation wages by choosing the occupations and destination countries they leave for, an important margin of adjustment that the administrative data allows me to explore. The spatial variation in destination countries within the Philippines is also critical for identification of the role of migrant demand conditions. For example, in a context where migration is predominantly to a single country, international migrant demand from the one country can still be critical, yet identifying its effects apart from aggregate time trends would not be feasible. Finally, the highly institutionalized migration system provides high quality administrative data, which is essential to construct precise measures of a variety of migration outcomes with geographical and temporal granularity.[3]

Related Literature

This study makes several contributions to the literature on the use of international migration as a shock coping device. The literature broadly finds that international remittance flows increase in response to negative weather and disaster shocks,[4] while evidence on the migration response is mixed,[5] reflecting the underlying heterogeneity across origin-destination pairs and different modes of migration (Munshi, 2003a; Mahajan and Yang, 2020a; Halliday, 2006; Beine and Parsons, 2017). I add to this literature by studying temporary international labor migration. While understudied, such migration is a common livelihood strategy across developing countries especially in South and Southeast Asia, and is expected to be more widespread in Sub-Saharan Africa (Adhikari et al., 2021). My key contributions are twofold. First, the focus on labor migration with the aid of administrative data allows me to provide novel evidence on how origin shocks lead to occupational and destination downgrading, likely due to falling reservation wages. Second, I document the role international migrant demand conditions play in mediating the migration response. My findings are related to Cinque and Reiners (2023), who show that an emigration ban in Indonesia negatively im-

[3]Data combining these migrant outcomes, migrant origins within countries, and temporal granularity is rare. Migrant-origin censuses and surveys tend to not include all of these variables, may miss migrants entirely, and surveys usually lack the necessary sample size for precise measurement at granular scales. Migrant-destination sources may include information on migrant earnings and occupations, but they do not allow for the study of country choice as they condition on showing up at the destination sources, capture a fraction of total migrants from the source country, and do not usually indicate migrant origin beyond the country.

[4]See Choi and Yang (2007); Blumenstock et al. (2016) for micro evidence. Yang (2008a); Mbaye and Drabo (2017) for cross-country evidence.

[5]Drabo and Mbaye (2015) and Marchiori et al. (2012) finds disasters and weather anomalies on average increase international migration. Cattaneo and Peri (2016) finds no average increase impact for natural disasters. Gröschl and Steinwachs (2017) also do not find an average impact, but notes that for middle income countries natural disasters drive out-migration, likely reflecting that middle-income countries are less financially constrained than the poorest, while less insured than the richest. Beine and Parsons (2017) documents a negative average effect, but points out migration to neighboring countries increase. Micro evidence on origin countries is also not equivocal, with, for example, Halliday (2006) and Yang (2008d) finding negative international migration responses to the 2001 earthquake in El Salvador while Giannelli and Canessa (2022) finding positive migration responses to flooding in Bangladesh. See Berlemann and Steinhardt (2017) and Cattaneo et al. (2019) for surveys of the literature.

pacted origin shock-coping capabilities. I complement this result by explicitly documenting international migration responses to shocks and demonstrating the importance of overseas migrant demand conditions as opposed to restrictions imposed by origin countries.

The focus on typhoons makes this work particularly relevant as the frequency and intensity of these extreme weather events are projected to increase due to climate change (IPCC, 2021; Kossin et al., 2017).[6] In closely related work, Mahajan and Yang (2020a) and Winter (2020) finds hurricanes and typhoons lead to increased international out-migration, primarily focusing on migration to the US.[7] Their results are driven by permanent migration through family visas. I complement these studies by focusing on temporary labor migration, which faces a different set of frictions and constraints than permanent migration to the US, making it a priori unclear whether similar results will hold in this context. My results suggest facilitating temporary labor migration can be a tool for dampening the impacts of typhoons.

More broadly, this study relates to the literature on international migration frictions (McKenzie et al., 2014; Bazzi, 2017; Bazzi et al., 2021b; Shrestha and Yang, 2019b; Beam et al., 2016; Fernando and Singh, 2023; Bryan et al., 2014; Naidu et al., 2023, 2016; Shrestha, 2019). The impacts documented in the paper are readily rationalized by presence of search frictions and congestion in the market for overseas contracts. Existence of such frictions have broader implications for the question "why don't individuals migrate more?" given the large documented income gains from international migration (Clemens, 2011a). While the number of global overseas contracts likely places a limit on how many individuals can migrate internationally through temporary labor migration (McKenzie et al., 2014), addressing the search frictions stemming from decentralized search may nevertheless increase migration through facilitating more migrant-job matches and induce more potential migrants to search for overseas jobs.

Finally, the paper speaks to two additional literatures. First, I offer additional insights regarding the well documented phenomenon of migrant downgrading in the labor literature (Dustmann et al., 2021, 2016; Eckstein and Weiss, 2004; Barsbai et al., 2019). I provide evidence that origin shocks can drive migrant downgrading when access to global labor markets are subject to search frictions. Further, I show that substituting to lower paying countries is another margin of downgrading for potential migrants along with occupational downgrading. Second, this study expands the recent empirical literature on how international migrant connections leads to propagation of shocks across countries (Chiswick and Hatton,

[6] Heavy storms account for nearly half of global economic damages and 17% of deaths due to natural disasters between 1998 and 2017 (CRED, 2017), and they have long-term impacts on the economic growth of affected countries (Hsiang and Jina, 2014)

[7] Focusing on Latin America, Hanson and McIntosh (2012) also finds natural disasters increase migration to United States, finding overall negative effects for all other countries.

2003; Gröger, 2021; Caballero et al., 2021; Khanna et al., 2022). I document an additional channel of propagation: destination country conditions can directly influence the capacity of migrant origin regions to absorb shocks through their effects on international migration.

The paper is structured as follows: Section 1.2 provides background information on temporary international labor migration. Section 1.3 presents the theoretical framework to guide the empirical analysis. Section 1.4 provides information on data sources and measurement. Section 1.5 presents results on the average effects of typhoons on migration outcomes. Section 1.6 presents the analysis of how migrant demand conditions coincident with typhoons mediate these effects. Section 1.7 presents the remittance analysis. Section 1.8 concludes.

1.2 Context: Temporary Labor Migration

Temporary international labor migration is exceedingly common, especially in lower-middle and low-income countries in South and Southeast Asia. It is a legal path to migration facilitated by signing an overseas work contract usually acquired through private intermediaries. The migrants travel alone and are required to return to their origin countries at the end of the pre-specified contract duration if they cannot renew their contracts. Destination countries are predominantly the Gulf States and other developed economies in Asia, with most migrants leaving for low- and semi-skilled occupations.

To demonstrate its scale, Appendix Figure A.1 plots the average annual temporary labor out-migrants between 2006-2016 for 12 countries with official statistics available.[8] Half the countries average over 500,000 new migrants annually, with the total across all countries surpassing 4.5 million. For Cambodia, Vietnam, and Myanmar, countries with fewest migrants, annual flows reach 100,000 by the end of the period (Appendix Figure A.1b). These large annual flows can correspond to significant shares of the working age population. Five countries send a percent or above of their entire working age population as temporary labor migrants every year, with the share reaching as high as 1.9% and 2.4% for Sri Lanka and Nepal.

Due to its large scale and potential importance for origin economies, facilitation and regulation of international labor migration is of policy concern to many countries. For example, of 70 developing countries with over a million population, 88% have a dedicated government agency for overseas employment, citizens abroad, or diaspora engagement (Khanna et al., 2022; United Nations, 2019). Temporary labor migration programs of many countries such

[8]The countries are Bangladesh, Cambodia, Egypt, India, Indonesia, Myanmar, Nepal, Pakistan, Philippines, Sri Lanka, Thailand, and Vietnam. The figures only include new migrant contracts, excluding contract renewals due to data limitations.

as India and Indonesia have been partially based on the Philippines example due to the perceived success of the Philippines' program (Theoharides, 2018a; Mishra and Rajan, 2010; Asis and Aguinas, 2012).

Evidence from the Philippines on the migration responses to natural disasters is likely to have broader applicability to countries that currently partake in similar international labor markets. Further, the analysis spans a period (2007-2016) where there were considerable cross-country competition in the international labor markets, increasing the relevance of the current results for the future where such cross-country competition across potential migrants is likely to persist.

1.2.1 Temporary Labor Migration in the Philippines

The Philippines is among the first countries to institutionally implement temporary international labor migration at scale with the 1974 Labor Code of the Philippines. Since then, temporary labor migration from the Philippines has risen significantly, with annual outflows increasing from 36,035 in 1975 to over a million in 2016 (IOM, 2013). 7.5% of Philippine households have an overseas labor migrant member and 2.2% of the population is composed of current migrants in 2015.[9]

Vast majority of migrant outflows from the Philippines are for temporary labor migration. Figure A.2 plots the annual migrant outflows from 2007 to 2016 for temporary labor migration and permanent migration as recorded by Commission on Overseas Filipinos (CFO).[10] Less than 10 percent of out migration is through permanent visas. The dominance of labor migration in terms of outflows, and greater attachment of temporary migrants to their communities and families back home, make it a particularly important migration channel to study as a shock-coping mechanism for origin regions.

Philippine labor migrants leave for a wide variety of occupations and destinations. For men, the top 3 occupations are production workers (11%), laborers/helpers (5.4%), and electrical wiremen (4.8%). For women, occupations are more concentrated, with domestic helpers (58%) being the most common occupation by a wide margin, followed by nurses (6.5%) and caregivers (5.5%). In terms of destinations, vast majority of migration is to Gulf countries and other developed nations in Asia (Appendix Table A.2). However, there is significant variation within the Philippines in which countries are prominent destinations, as documented in Appendix Figure A.7. For example, from 2006 to 2017, the share of migrants

[9] Calculated from 2015 Census microdata.

[10] The other legal alternative to temporary labor migration is to emigrate with a permanent visa, over 90% of which is through family-sponsored, as opposed to employer-sponsored, visas. Filipinos who hold a permanent immigrant visas must legally register with the CFO before departing.

leaving for Hong Kong is 25%-40% for many municipalities in northern Luzon, while the average municipality in the Southern island of Mindanao only sends 3.8% of migrants to Hong Kong. Conversely, migration to Kuwait is heavily concentrated in the south of the country. These persistent differences in migrant shares generate substantial heterogeneity in the incidence of destination country shocks within Philippines.

The wages of migrants are regulated through a host of regulations by the Philippines government, bilateral agreements between the Philippines and destination countries, and labor market regulations of destination countries. Work contracts that do not conform to the relevant regulations are not approved by the Philippines government. Chief among such regulations are minimum wages, such as the $400 minimum wage for all Filipino migrants leaving for the domestic service occupations, enacted in 2006. Philippine Overseas Labor Offices is tasked with ensuring any verified contract is in accordance with both domestic and overseas regulations, along with making sure the contract wages are in line with prevailing market wages of the host country for the occupation at hand (McKenzie et al., 2014). Therefore, even for higher paying occupations and countries, where domestic or host minimum wages may not bind, there are strict limits on how low contract wages can be. Overall, the regulatory setting makes it unlikely that overseas wages of Filipino's can downward adjust easily in the short run in response to domestic migrant supply shocks. As shown in McKenzie et al. (2014) (replicated and extended in Appendix Section A.5.4) these regulations act as a binding minimum wage for Filipino migrants that lead to excess supply of migrants at prevailing wages which tend to offer substantial premiums over earnings at home.

Temporary labor migration is primarily facilitated by licensed private intermediaries. Potential migrants find overseas contracts at home through these intermediaries to secure their exit visa. Recruitment agencies provide many services that includes matching with employers, filling information gaps potential migrants may have, and logistical support to navigate the legal requirements of the recruitment and migration process. Potential migrants can connect with intermediaries directly via office visits in cities or online. It is also common practice for recruitment agencies to work with informal brokers who have access to more remote areas or organize job fairs in different parts of the country.

Searching for overseas jobs is a decentralized process between potential migrants and multitude of private recruitment agencies, of which there were over a thousand in the analysis period.[11] This process can be subject to significant search frictions and possibility of failed search. According to a 2022 survey, 17% of adults in the Philippines had aspirations to work

[11] There were licensed 1163 recruitment agencies in May 2022.

abroad and 7% actively searching, a fraction much larger than the annual migration rate.[12] Additionally, Beam et al. (2016) reports results from an experiment in the Philippines where a migration facilitation intervention caused a sizable increase in number of people searching for overseas work, yet substantially smaller increase on whether they were able to secure an overseas contract. Of course, failed search on the migrant side can reflect search frictions or excess supply of migrants relative to jobs. However, there are reports of unfilled vacancies even in the presence of migrants searching for work, consistent with search frictions impeding matches from taking place. According to one report "... there were contracts that were not supplied with workers: 162,823 in 2006, 228,254 in 2007; and, 228,282 in 2008,"[13] suggesting unfilled overseas vacancies are common.

1.3 Theoretical Framework

To guide the empirical analysis, this section provides a model of international migration subject to search frictions, combining elements of canonical Diamond-Mortensen-Pissarides and partial equilibrium search frameworks. Consistent with the empirical setting, overseas wages can offer significant premium over origin earnings, and migration is conditional on securing a foreign contract before leaving. This process is subject to search frictions. I start with a model of random search. In this setup, I explore the impacts of a typhoon shock that decreases home earnings. Returns to migration increase and reservation wages fall, therefore increasing out-migration and lowering migrant wages. Better migrant demand conditions abroad at the time of the shock (which decreases the cost of recruitment agencies for finding overseas vacancies) leads to a larger migration response. I then modify the model by allowing migrants to direct their search across a *high* and a *low* paying overseas market, given potential migrants can direct their search somewhat. Key conclusions remain the same. Additional to the random search model, migrants respond to the shock by shifting search into lower paying labor markets with higher likelihood of securing a job, reinforcing the drop in migrant wages in response to typhoons. I conclude by discussing robustness of the models predictions to maintained assumptions. The details of the analysis and derivations are presented in Appendix Section A.3.

[12] https://www.sws.org.ph/swsmain/artcldisppage/?artcsyscode=ART-20230419092946&mc_cid=3706e251e0&mc_eid=1eeee26a57

[13] https://news.abs-cbn.com/pinoy-migration/11/10/09/rehired-land-based-ofws-boost-worker-deployment-last-5-years

1.3.1 Environment

Basics. The model is in discrete time. There is a large number of homogeneous, infinitely lived, and risk neutral individuals who apply discount factor β to future utility. Individuals earn y_o per-period at home. They choose between staying at home or searching for an overseas work contract. If they search, they incur the search cost c.

Search. Finding a contract in a foreign labor market is subject to search frictions. If individuals choose to search, they match with an overseas contract with probability $q(s, v)$ which is decreasing in number of searchers s (due to congestion) and increasing in job availability v (for vacancy). Matches are realized with a Cobb-Douglas matching function: $m(v, s) = s^{\alpha}v^{1-\alpha}$.[14] The per-period probability an individual finds a contract is given by the usual matches over mass of searchers: $q(s, v) = \frac{s}{v}^{-\alpha}$. If matched, the contract offers utility $w_t \in [\underline{W}, \overline{W}]$, drawn from the distribution $F_w(.)$. The individual can choose to accept the offer or remain home after w_t is revealed. While w_t can be thought of as comprised of both the pecuniary and non-pecuniary costs and benefits of a given contract, I assume expected wage is increasing in w_t and will refer to w_t as wages for the rest of the section.[15]

Individuals' Problem. The expected lifetime utility associated with searching for overseas contract (V^s) and not searching (V^o) are:

$$V^o = \frac{y_o}{1-\beta} \qquad V^s = y_o - c + \beta q(s, v) \int_{\underline{W}}^{\overline{W}} \max\{w_t, V^s\} dF_w + \beta(1 - q(s, v))V^s \qquad (1.1)$$

Conditional on choosing to search for an overseas contract, this setup implies a reservation wage policy where migrants accept an offer if $w_t \geq w^* = V^s$. Further, in equilibrium, homogeneous workers are indifferent between searching for an overseas contract or not searching: $V^s = V^o$.[16] Two conclusions follow. First, the discounted expected gain from searching equals the search cost in equilibrium:

[14] As shown in Appendix Section A.3, all results in this section goes through with a general CRS matching function. I use Cobb-Douglas in the body for ease of exposition.

[15] Note that w_t denotes the full value function associated with an offer, encoding information about its duration as well. For example, an overseas contract that offer per period utility x for two periods would have $w_t = \beta x + \beta^2 x + \beta^3 V^s$, capturing that after two periods abroad the migrant returns home and is back to searching. While the length of a contract would effect its associated value w_t, over 75% of contracts in the data are for two years. Given low variation in contract length, I assume variation in w_t is primarily from contract wages.

[16] I assume that the mass of individuals are large enough that a corner solution where everybody searches for foreign jobs and still $V^s > V^o$ does not emerge.

$$\beta \underbrace{q(s,v)}_{\mathbb{P}(offer)} \underbrace{\bar{F}_w(w^*)}_{\mathbb{P}(accept|offer)} \underbrace{(\mathbb{E}[w_t|w_t > w^*] - w^*)}_{\equiv\Delta(w^*)} = c \tag{1.2}$$

where $\bar{F}_w(w^*) = \mathbb{P}[w_t > w^*]$ and $\Delta(w^*)$ is the gap between expected foreign utility beyond the reservation wage in case an individual matches with and accepts a foreign contract. Second, the reservation value $w^* = V^o = \frac{y_o}{1-\beta}$ in equilibrium.

Recruitment Agencies. Access to overseas contracts are mediated by a representative overseas recruitment agency. The agency earns fixed revenue p for every accepted offer, and faces convex vacancy costs. The per period recruitment agency profit is $\Pi = pm - v^\rho$ where $\rho > 1$ mediates the convexity of the costs and m is the total migration that takes place in the period. The increasing marginal costs capture the fact that while initial vacancies can be essentially very low cost as recruitment agencies have established contacts abroad, finding additional vacancies would require higher effort cost due to the limited number of global vacancies available in a period. The steepness of the cost curve depends on the migrant demand conditions, which decreases the cost for the firm to secure additional jobs. The firm picks per-period v to maximize per period profit Π.

1.3.2 Equilibrium and the Impact of Origin Shock

The equilibrium number of migrants m and average migrant wages $\bar{w}$ are given by:

$$m = \left(\frac{p}{\rho}\right)^{\frac{1}{\rho-1}} \left(\frac{\beta\Delta(w^*)}{c}\right)^{\frac{1-\alpha}{\alpha}\frac{\rho}{\rho-1}} \bar{F}_w(w^*)^{\frac{1}{\alpha}\frac{\rho}{\rho-1}} \qquad \bar{w} = \mathbb{E}[w_t|w_t > \frac{y_o}{1-\beta}] \tag{1.3}$$

Intuitively, migration is increasing in foreign wages, as reflected through the expected gap between the average accepted overseas wage versus reservation wages ($\Delta(w^*)$) and the likelihood that a given contract would be above the reservation wage ($\bar{F}_w(w^*)$). Migration is falling in cost of searching c and the steepness of the cost curve of recruitment agencies ρ.

Consider the impact of a negative origin shock which decreases home earnings ($dy_o < 0$) which uniformly increases the returns to migration to each foreign labor market:[17]

[17]For simplicity, I assume the change are permanent, though transitory changes in y_o would lead to same qualitative conclusions.

$$\frac{\partial \ln m}{-\partial y_o} = \frac{\rho}{\rho - 1}\frac{1}{1-\beta}\Bigg[\underbrace{\frac{\alpha^{-1}}{\Delta(w^*)}}_{\uparrow search} - \underbrace{\frac{1}{\Delta(w^*)}}_{\downarrow \mathbb{P}(offer)} + \underbrace{\frac{f(w^*)}{\bar{F_w}(w^*)}}_{\uparrow \mathbb{P}(accept|offer)}\Bigg] > 0 \tag{1.4}$$

$$\frac{\partial \ln \bar{w}}{-\partial y_o} = \frac{1}{1-\beta}\frac{f(w^*)}{\bar{F_w}(w^*)}\left(\frac{w^*}{\bar{w}} - 1\right) < 0 \tag{1.5}$$

Equation 1.4 captures that an increased returns to migration increases both the number of searchers and the share of searchers who accept their overseas offer due to lower reservation wages. During better demand conditions (lower ρ) migration esponse is larger as recruitment agencies are better able to increase vacancies in response to increased supply of potential migrants. The drop in reservation lead to a drop in average migrant wages, as shown in 1.5. These expressions imply the following results.

Result 1 (Migration Response). *If typhoons primarily decrease home utility, they lead to increased out-migration and lower average migrant wages due to lower reservation wages.*

Result 2 (Role of Migrant Demand Conditions). *Better migrant demand conditions during a shock (lower ρ) magnify the migration response without affecting average migrant wage response. Therefore, migration elasticity of migrant wages $\frac{d\ln\bar{w}}{d\ln m}$, i.e. the percent drop in wages per one percent increase in migration, fall in magnitude.*

In reality, origin shocks can both decrease origin utility and increase barriers to migration. Appendix Section A.3, considers this case by allowing the shock increases the fixed search costs ($dc > 0$). This is meant to capture variety of unmodeled channels a natural disaster can impede migration, such as loss of necessary documents due to damages,[18] increases in the price of migration services and assistance in response to increased demand, destruction of infrastructure making it harder to access recruitment services, or increased inability to pay the fixed cost of migration due to asset/wealth losses or disruptions to migration financing.

Result 3 (Increased Barriers to Migration). *If typhoons both decrease home utility and increases barriers to international migration, they have an ambiguous affect on out-migration based on which force dominates. They still decrease reservation wages and average*

[18]See, for example, https://www.thenationalnews.com/world/filipinos-seek-middle-east-jobs-to-rebuild-lives-after-haiyan-1.260462 which reports in the context of the 2013 Typhoon Haiyan that "[m]any lost their passports, birth certificates and certificates of employment in the storm surge that followed the typhoon. At the Tacloban job fair, half of those applying for overseas jobs did not qualify because of a lack of documents."

migrant wages fall.

The sign of the migration response to shocks is therefore informative about whether increased barriers to migration due to disasters dominates the incentives caused by increased returns to migration.

1.3.3 Extension: Multiple Foreign Labor Markets

The analysis so far assumes search for overseas contracts is entirely random. In reality, migrants can direct their search across foreign markets by, for example, approaching recruitment agencies that specialize in certain destination countries or occupations. To assess whether directed search changes predictions regarding migrant wage drop following typhoons, I extend the model by allowing individuals to direct their search between two labor markets, denoted by h for *high* wage and l for *low* wage.[19] The individuals still face an exogenous distribution of wages in each market, with the distribution of wages in h greater than in l in hazard rate ordering, i.e. $\frac{f_w^h(w_t)}{\bar{F}_w^h(w_t)} \leq \frac{f_w^l(w_t)}{\bar{F}_w^l(w_t)}$.[20] The overseas markets are otherwise identical.[21]

As before, in equilibrium, individuals are indifferent between searching in either of the overseas markets or not searching. Therefore, reservation wage for both markets are equal to $w_h^* = w_l^* = \frac{y_o}{1-\beta}$ and condition 1.2 holds for both overseas markets. This implies:

$$\frac{q(s_h, v_h)}{q(s_l, v_l)} = \frac{\bar{F}_l(w^*)\Delta_l(w^*)}{\bar{F}_h(w^*)\Delta_h(w^*)} \leq 1 \tag{1.6}$$

Probability of matching with a contract in the high wage market is lower than it is for the low wage market in equilibrium. Intuitively, the high wage overseas market attracts relatively more individuals to search for contracts, pushing down the likelihood of a match for any individual searcher.

[19] Generalizing to N locations that are ordered does not change the conclusions of this section, as shown in Appendix Section A.3.

[20] Hazard rate ordering is a way to stochastically order two random variables, i.e. define one variable to be "bigger" than the other. Intuitively it captures that wages in market h are higher than wages in market l. Hazard rate ordering is stronger than the usual first-order stochastic dominance ($\bar{F}_w^h(w_t) \geq \bar{F}_w^l(w_t)$) and weaker than likelihood ratio ordering. (Ross 1983)

[21] The results in this extension requires an additional assumption on the match function: $\eta(\theta)$ is weakly decreasing in θ, where $\eta(\theta) = \frac{d \ln q(\theta)}{d \ln \theta} \in (-1, 0)$ is the elasticity of the job finding probability with regards to the market tightness and $\theta = \frac{v}{s}$ is the market tightness (Appendix Section A.3). Cobb-Douglas, along with other common matching functions in the search literature like CES with gross complements and urn-ball matching function satisfies this condition. The condition implies that in more congested markets a percent change in the relative number of searchers lead to weakly lower percent change in the job finding probability than in less congested markets. Failing this condition is not sufficient for the results of this section to not hold.

The mean migrant wages are now given by the weighted average $\bar{w} = \pi_h \bar{w}_h + \pi_l \bar{w}_l$ where π_i is the share of migrants going to market i and $\bar{w}_i = \mathbb{E}_i[w_t | w_t > \frac{y_o}{1-\beta}]$ is the mean wage of migrants going to i. In response to an origin shock $dy_o < 0$, the average migrant wage response is a function of both the changes in mean migrant wages in each overseas market and the changes in the share of migrants going to each market:

$$\frac{\partial \bar{w}}{-\partial y_o} = \frac{\partial \pi_h}{-\partial y_o}(\bar{w}_h - \bar{w}_l) + \frac{\partial \bar{w}_h}{-\partial y_o}\pi_h + \frac{\partial \bar{w}_l}{-\partial y_o}\pi_l \tag{1.7}$$

The average wage within markets fall ($d\frac{\partial \bar{w}_h}{-\partial y_o} < 0$ and $d\frac{\partial \bar{w}_l}{-\partial y_o} < 0$) given the drop in reservation wages, as in the case of fully random search. However, the total effect on the average migrant wages is also determined by the change in the share of migrants going to the high paying market. Appendix Section A.3 shows $\frac{\partial \pi_h}{-\partial y_o} < 0$ as a uniform drop in home utility increases the relative attractiveness of the low wage market with its higher likelihood of securing a contract.

Result 4 (Heterogeneous Foreign Markets). *If individuals can direct their search across heterogeneous (in terms of wages) overseas labor markets, a negative shock to origin income increases the share of potential migrants searching in and migrants going to the lower paying overseas market.*[22] *Therefore, average migrant wages fall both due to lower reservation wages of migrants and increased share of migrants going to lower paying overseas markets.*

Therefore, both the random and directed models of search leads to consistent predictions regarding the average migrant wage response to typhoons. The empirical analysis will not be able to differentiate between whether changes in the share of migrants going to high/low paying countries/occupations are driven by migrants directing search towards these jobs or searching more generally but accepting lower paying "draws", with reality likely reflecting a combination of both channels.

1.3.4 Discussion of Assumptions

Before the empirical analysis, I address the plausibility and importance of the simplifying assumptions of the model. That home earnings are not influenced by the migration rate is plausible given typhoon-induced migration make up a very small fraction of origin working age population. Moreover, this assumption can be relaxed without qualitative changes to the

[22] While first-order stochastic dominance guarantees that share of individuals searching in the lower paying overseas market increases, hazard rate ordering ensures that share of migrants going to lower paying market increases as well. See Appendix Section A.3.

predictions of the model. Increased origin wages due to out-migration would dampen, but not reverse the migration response. That foreign wages are exogenous from the perspective of the origin communities is also realistic. Migrants from any origin municipality makes up a small fraction of destination labor force, and migrant wages are regulated by both the Philippines and destination countries. Further, Appendix Section A.5.3 shows that migrant wages do not vary by origin within Philippines after conditioning on destination country and occupation, alleviating worries about localized migrant wage responses. Finally, in the model, foreign demand conditions only affect the marginal cost of recruitment agencies vacancy posting costs, but do not change the overseas wage distribution. This is consistent with the findings of McKenzie et al. (2014) in the Philippines, Bossavie et al. (2021a) in Bangladesh, and Appendix Section A.5.4 of this paper where migrant quantities respond to GDP shocks in destination countries without a change in migrant wages.

A key simplification of the model is the homogeneity of individuals. This choice is driven by the goal of showing how migrant wages can respond in the absence of any selection. However, if potential migrants have heterogeneous home utility or productivity, the marginal migrant would have higher reservation wages, making the overall migrant wage response ambiguous due to competing forces of positive selection and falling reservation wages. Which of these forces dominate is an empirical question, but the interpretation of a negative empirical wage result is unchanged.

Further, incidence of typhoons or the ability to respond may vary across different groups, leading to an additional source of selection. For example, in the presence of worsening liquidity constraints, only richer and more educated households may be able to respond, putting upward pressure on migrant cohort wages. Conversely, richer households may be better insured or adapted against shocks, shifting the migrant composition towards poorer and lower educated populations. I explore these issues in the empirical section of the paper. I present evidence that migrant cohorts following typhoons are more educated. Decreasing wages in the face of increasing migrant education would further underline the importance of the reservation wage impacts underlined in the model.

1.4 Data, Measurement, and Summary Statistics

This section discusses the main data sources and measurement of variables of interest. For brevity, I discuss measurement at the municipality-year level. Variables with different spatial or temporal dimensions (i.e. province or bi-quarterly) are created analogously. More details on data construction and definition are provided in Appendix Section A.2.

1.4.1 Migration: Administrative Contract Data

The key migration dataset is from the administrative database of the Philippine Overseas Employment Administration (POEA). Before migrating, all temporary labor migrants are required to visit the POEA to have their contract approved and to receive exit clearance. This results in POEA maintaining a dataset of all new migrant contracts from the Philippines.

I have access to the dataset of all land-based contracts leaving the Philippines from 1992 to 2016. The data includes information on sex, date of birth, contract occupation, destination country, and the salary of each migrant contract. For two periods in the data I also observe information about migrant's home address: 1992-1997 and 2007-2016.[23] This is critical for my research design as the unit of analysis is Philippine municipalities. I use the 1992-1997 period as the *baseline period*, which I use to construct baseline variables. 2007 to 2016 is the *analysis period*, throughout which I observe 4 million migrant contracts, with over 90% containing origin municipality information.

The main outcomes of interest are the migration rates and the wages of migrants leaving each municipality. I construct the migration rate by dividing the number of migrants by interpolated population calculated from the 2007, 2010, and 2015 Population Census. To construct the wage measure, I convert all salary information to one-year equivalent real 2010 Philippine Pesos (Phps). The wage measures are then constructed as the mean (along with 25th percentile, median, and 75th percentile) of the wage distribution of migrants leaving a municipality in a year.

Finally, I measure the share of migrants going to high versus low paying countries and occupations in a given year. I use the baseline period data to group countries and occupations into quartiles based on their wage levels. To do so, I regress log individual wages jointly on occupation and destination country fixed effects. I then collect these estimated fixed effects and use the empirical bayes shrinkage estimator of Morris (1983) to account for noise in the estimation leading to possible bias. I then group countries and destinations in quartiles based on the value of the fixed effects, with an equal number of occupations and countries in each quartile. Table 1.1 shows the top two destinations and occupations from each quartile. With these groupings, I calculate the share of migrants going to each occupation and destination quartile for each municipality-year.

[23]Starting from 2010, the POEA data includes migrants home municipality and province. For the previous years, I rely on a matched dataset between the POEA database and the Overseas Worker Welfare Administration (OWWA) database, which was the government agency responsible for the well being of overseas workers and their families. The OWWA database includes information about migrant's home address. The matched database is created through a fuzzy matching algorithm that uses the first name, middle, name, last name, date of birth, destination country, sex, and year of departure of the migrants (Theoharides, 2018a). A 95% match rate is achieved.

Table 1.1: Top Countries and Occupations in Each Wage Quartile (2007-2016)

	Country	**Count**	**Occupation**	**Count**
1st Quartile	Saudi Arabia	1,404,274	Domestic Helper and Related	1,387,383
	UAE	542,319	Laborer	338,188
2nd Quartile	Libya	20,235	Plumber, Welder, and Related	200,810
	Cyprus	14,543	Bricklayer, Carpenter, and Related	116,180
3rd Quartile	Taiwan	332,949	Clerical and Related	80,498
	Israel	15,409	Material-Handling Equipment Handlers	77,663
4th Quartile	Hong Kong	264,421	Medical, Dental, and Related	192,650
	Japan	57,167	Engineers, Architects and Related	196,991

Notes: Top two destinations and occupations within each quartile grouping. Migrant counts are from 2007-2016.

1.4.2 Typhoon Exposure Measurement

Typhoons are a potentially highly destructive form of tropical cyclone that form in the Northwestern Pacific basin. Philippines is among the most typhoon exposed countries in the world, with approximately 20 tropical cyclones entering the region surrounding the country (named Philippine area of responsibility) annually. A subset of these cyclones reach typhoon scale winds and make landfall in the Philippines every year, causing considerable damages and welfare loss (Franklin and Labonne, 2019).

Typhoons vary considerably in their intensity (particularly wind speed), exact location, impact area, and how populated the affected areas are. The stronger the wind speeds and the more populated the areas they impact, the higher the economic (and human) damages usually are. Accordingly, I construct a typhoon exposure index broadly following Mahajan and Yang (2020a) that accounts for these features. The meteorological nature of the index ensures that it is not prone to error or bias due to misreporting. Below I describe the steps to create the municipality-year typhoon exposure index broadly, relegating equations and details to Appendix Section A.2.1:

1. Use Joint Typhoon Warning Center (JTWC) best-track data to estimate the maximum wind speed that prevailed in every 30 arc-second grid cell for each tropical cyclone.

2. For grid-cells reaching wind speeds above tropical storm speed winds (34 knots), normalize max wind by subtracting this threshold (34 knots) and dividing by the maximum wind speeds observed in the data.

3. Aggregate to municipality-year level by taking a weighted sum of the normalized maximum wind speed across grid-cells and storms that are within a municipality-year. Cells

are weighted by the population residing in the cell as reported by the Socioeconomic Data and Applications Center gridded population of the world for the year 2000.

4. Normalize the sum by dividing it by the municipality population to ensure larger values are not driven by size of municipality.

The resulting index can be interpreted as the intensity-weighted per capita typhoon exposure in a municipality-year, where intensity is both driven by the number of storms that hit the municipality in a year, and by the wind speed of each storm. Throughout the paper, I standardize the index to have mean 0 and standard deviation 1 for ease of interpretation.

Figure 1.1 visualizes the municipality-year level typhoon index for three consecutive years in the analysis period: 2011, 2012, and 2013. The left panel shows the path of all tropical depression level storms that passed through the Philippines in a given year, along with the maximum predicted wind speeds across the year in each pixel. The right panel shows the final typhoon exposure index T_{mt} at the municipality level. The figure underlines how common typhoons are in the Philippines, and also visualizes the variability in the exact location and intensity of the typhoons, which is the variation I exploit in my empirical analysis.

1.4.3 Other Data Sources

I use a variety of other data sources throughout the paper, including the Population Censuses, Family income and Expenditure Surveys, Surveys on Overseas Filipinos, government typhoon damage estimates, and nighttime light intensity. I discuss the sources and resultant variables in the relevant sections of the paper.

1.4.4 Summary Statistics

Appendix Table A.1 presents summary statistics over the analysis period 2007-2016. The average municipality-level annual migration rate is 0.34 percent and ranges from zero to 4.2 percent of population. Average yearly migrant cohort wages are \$5371 and vast majority of contracts are for approximately 2 years. Migrants primarily leave for low-paying countries and occupations, with average share leaving for the lowest wage quartile occupations at 67% and lowest wage quartile destinations at 77%. Finally, migrants tend to be more educated than the general population. The average educational attainment of migrant stock is 85% for high school completion and 39% for college completion, while corresponding values for the working age population is 57% and 14%.

Migrant Wage Variance Decomposition. In Appendix Section A.5.3, I undertake a variance decomposition exercise for contract wages. Destination country accounts for 38%,

Figure 1.1: Maximum Wind Speeds and Typhoon Exposure Index From 2011 to 2013

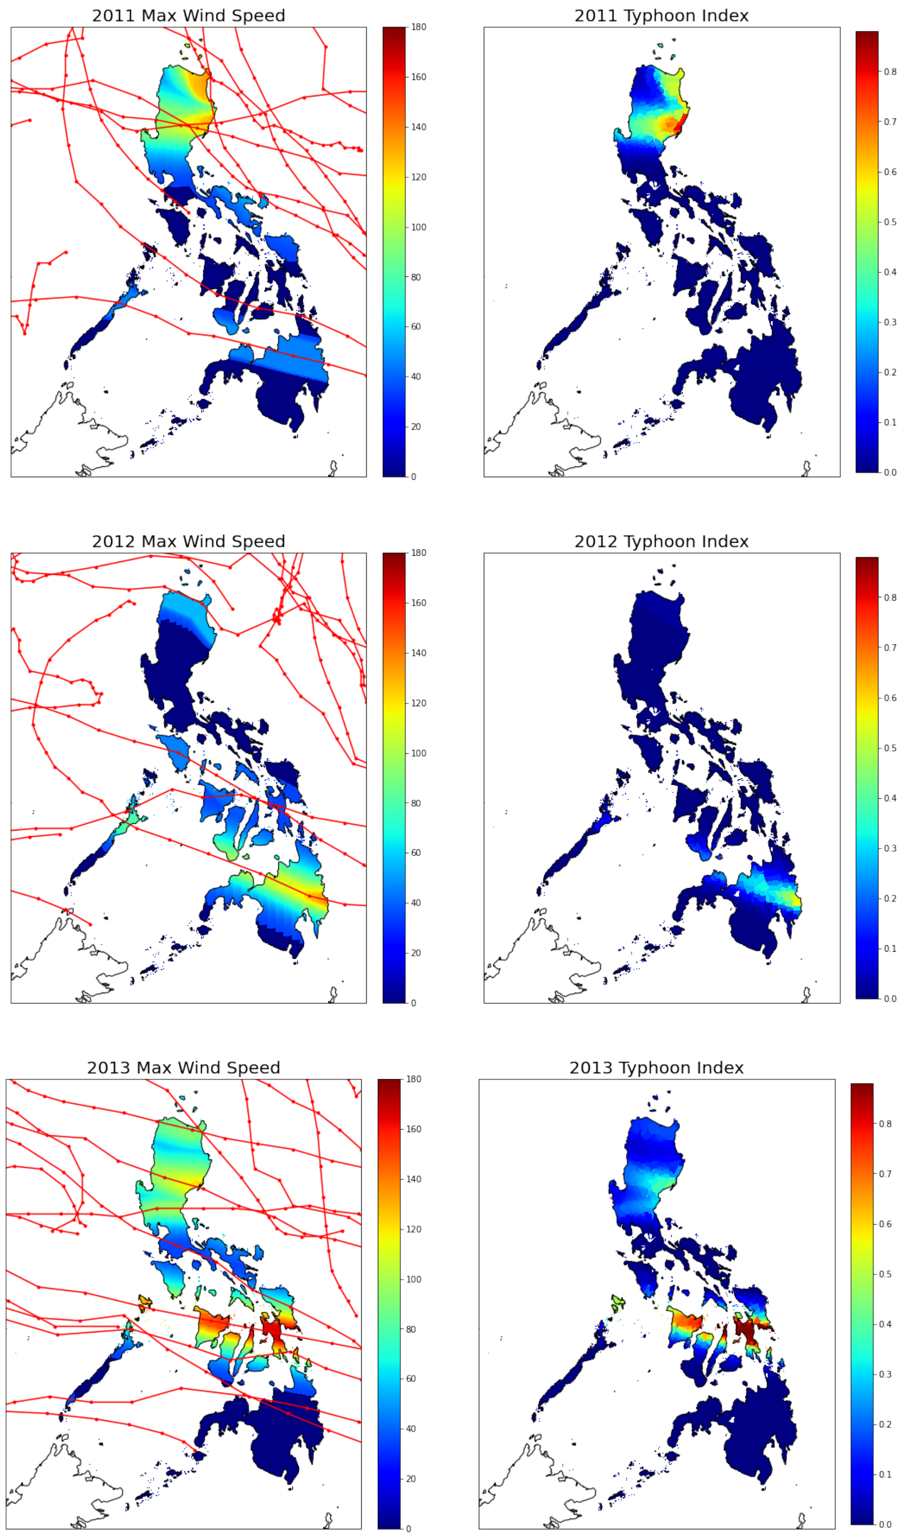

Notes: Left panel shows the path of each tropical depression passing the Philippines (red lines) and presents the maximum wind speed (if $\geq$ 34 knots) that prevailed in each 30 arc-second grid cell. The right panel shows the resulting municipality-year level typhoon exposure measure T_{mt}

occupation 32%, and the demographics (sex and age) 2% of the variation in the wages. Around 28% of variance remains unexplained. With more granular occupation-by-country cells, 76% of variation is explained by occupation and country, with 22% of the variation still unexplained. Overall, occupation and destination choices of the migrants explain a substantial portion of the wage variation.[24] This finding anticipates that occupation and destination choices of migrants will be an important margin driving the changes in migrant cohort wages in response to typhoons. Further, migrant wages do not vary by municipality of origin after conditioning on destination country, occupation, and demographics (Appendix Table A.17). Therefore, wages of contracts should not be directly responding to events local to a municipality such as a natural disaster.

1.5 Migration Responses to Typhoons

1.5.1 Validating the Typhoon Exposure Index

I validate the constructed typhoon exposure index by showing that it predicts physical and economic damages. To do so, I first obtained data on province level typhoon damages and casualties from the Philippines government. As detailed in Appendix Section A.5.1, the exposure index is a strong predictor of province-year level number of casualties, number of people affected, and pecuniary damage estimates due to typhoons. Further, a one standard deviation typhoon exposure in a given province-quarter leads to approximately a 2% drop in nightlight intensity. These results provide strong evidence that the typhoon exposure measure tracks typhoon-related damages and destruction. I interpret these results as a drop in home utility in typhoon affected regions, which increases the returns to migration. The rest of the section focuses on migration responses to typhoons.

1.5.2 Empirical Approach

To estimate the causal effects of typhoon exposure, I use a fixed effect strategy that exploits the exogeneity of the exact location, intensity, and timing of typhoons. I primarily employ a summary specification throughout the paper that summarizes the short run (0 to 1 years) and medium run (2 to 3 years) effects of typhoons. Additionally, I show results from an event study specification that allows me to assess dynamics more granularly and check for

[24] Appendix Section A.5.3 presents additional details and full set of results on the decomposition exercise. The decomposition results here refers to Panels B and C of Appendix Table A.16 and splits the covariance terms equally across the groups of regressors.

pre-trends. The main migration results are at the level of 1597 Philippine municipalities for 2007-2016.[25] The empirical specifications employed throughout the section are as follows.

Summary Specification. To summarize the effects of typhoons on an outcome of interest in the short- and the medium-run, I estimate the following regression as my main specification:

$$y_{mt} = \beta_{SR} T_{m,(t,t-1)} + \beta_{MR} T_{m,(t-2,t-3)} + \gamma_m + \gamma_{r(m),t} + \epsilon_{mt} \tag{1.8}$$

where y_{mt} is the outcome of interest for municipality m in year t. $T_{m,(t,t-1)}$ is the average of municipality m's typhoon exposure in the past two years ($T_{m,t}$ and $T_{m,t-1}$). Similarly, $T_{m,(t-2,t-3)}$ is the average of municipality m's typhoon exposure two to three years ago ($T_{m,t-2}$ and $T_{m,t-3}$). γ_m is the municipality fixed effect that controls for any time-invariant characteristics of municipalities. $\gamma_{r(m),t}$ is the island-group by year fixed effect that controls flexibly for any aggregate shocks or differential trends within Filipino island groups over time.[26] I calculate and present two sets of standard errors. First, I cluster the standard errors at the province level, which is a more aggregate administrative unit dividing the country into 79 provinces. Second, given the spatially correlated nature of both typhoons and many of my outcomes of interest, I present spatially clustered standard errors following Conley (1999), allowing for up to 200 kilometers around the centroid of the municipality and for auto correlation of order 10 years. For analyses where the outcome is a characteristic of migrant cohorts (for example average wage or share female), I weight the observations by the number of migrants used to calculate the characteristic y_{pt}, i.e. the cell size.[27]

The coefficients of interest are β_{SR} and β_{MR}. The coefficient β_{SR} can be interpreted as the effect of one standard deviation increase in the average typhoon exposure across the current and the previous year, which I refer to as the short-run effect. Similarly, β_{MR} can be interpreted as the effect of one standard deviation increase in the average typhoon exposure across two and three years ago, which I refer to as the medium-run effect.

Event Study Specification. I further employ an event study specification for main outcomes to document granular temporal dynamics and assess pre-trends. The estimating

[25] For individual migrant level regressions, the level of variation for the typhoon exposure is still the municipality.

[26] Island groups are the the largest administrative unit in the Philippines, dividing the country into three administrative regions: Luzon, Visayas, and Mindanao.

[27] This approach is common in migration literature (Bertoli et al., 2017; Borjas, 2003; Mishra, 2007) to ensure that estimates are not driven by severely noisy observations with few migrants. Further, weighting by number of migrants making up each cell allows me to recover estimates that are consistent with the migrant-level regressions I employ later in the paper.

equation is:

$$y_{mt} = \sum_{\tau=-K,\tau\neq-1}^{\tau=T} \delta_\tau T_{m,t-\tau} + \gamma_m + \gamma_{r(m),t} + \epsilon_{mt} \tag{1.9}$$

where $T_{m,t-\tau}$ is the typhoon exposure index of municipality m at τ periods before t, and y_{mt}, γ_m, $\gamma_{r(m),t}$ are the same as the summary specification above.

Identification and Discussion of Specification. Identification stems from the natural random variation in exact location, severity, and timing of typhoons. Specifically, the key identifying assumption is that the deviations from a municipality's average typhoon exposure in a given year is uncorrelated with other determinants of migration outcomes, conditional on municipality and year fixed-effects. This is reasonable given, while regions differ in their overall exposure to typhoons, predicting the intensity of a single typhoon season with geographical granularity, which depends on the timing, severity, and paths of individual typhoons, is difficult (Deryugina, 2017).

Callaway et al. (2021) studies the two-way fixed effects (TWFE) estimator with continuous treatment. For the estimates to be interpreted as a weighted average of causal responses to an incremental change in typhoon exposure, one needs to rule out possible selection bias stemming from units "selecting into" treatment based on treatment effects. This is reasonable in my setting, as the randomness of exact location, timing, and intensity of typhoons rules out the possibility of regions "selecting into" a particular treatment dose (deviation from average typhoon exposure) in a given period.

More broadly, the estimating equations 1.8 and 1.9 fall under the umbrella of TWFE estimators, applied in a setting with variation in timing of a continuous treatment, multiple treatments per unit, and with less than 2.5% of units never treated. A recent and growing literature raises concerns that, with staggered treatment timing, the presence of treatment heterogeneity across units or over time can contaminate the difference-in-difference estimates and the leads and lags in event studies (Roth et al., 2023). While this literature has provided robust estimands under a variety of research contexts (such as binary treatment or pure control group, see Sun and Abraham (2021a) and Callaway and Sant'Anna (2021)), to the best of my knowledge, there is no robust estimator in the literature corresponding to all the characteristics of my setting. Note that, in my setting, the randomness of typhoons implies that earlier or later treated municipalities are not selectively different, and therefore the causal responses to treatment are equal in expectation across different treatment timing-groups. This reduces concerns of bias stemming from comparisons across timing-groups with heterogeneous treatment effect paths.[28]

To further assess the possibility of bias in TWFE estimates, I follow the stacked-by-event design of Cengiz et al. (2019) using a binarized version of the typhoon exposure index in Appendix Section A.5.2. I check, with the binarized exposure, whether the TWFE estimates and stacked-by-event design estimates differ significantly, which could indicate the TWFE estimates are biased in this context. The stacked-by-event design ensures that the estimates are not influenced by "bad comparisons" by constructing a control group for each treatment-cohort that have not been treated for a window around the treatment-cohort of interest. I focus on a window of three years before and after each 6-month period (cohort). I show that for binary typhoon measures constructed as high typhoon exposure, the TWFE and stacked-by-event estimates track each other closely, alleviating concerns about biased TWFE estimates due to treatment effect heterogeneity. I proceed with using the continuous typhoon-exposure measure in the paper.

1.5.3 Typhoons Increase Migration and Decrease New Migrants' Wages

I begin by documenting that typhoons lead to an increase in international labor migration from affected municipalities. Column 1 of Table 1.2 presents results of the main summary specification. One SD average hurricane exposure in the past two years increases the migration rate by 1.3 migrants per 10,000. This increase persists in the medium-run, with an increase of 1.5 migrants per 10,000. These are economically meaningful effects corresponding to 3.7% and 4.3% of the mean migration rates. To assess the dynamics more granularly, Figure 1.2a presents results from a quarterly event study specification. Effects of a typhoons start manifesting in two quarters, is persistent for 15 quarters, and dissipates afterwards. There is no evidence of a short-run increase in migration being offset by a following drop in migration for up to 4 years. Overall, these results show that international labor migration is used as a shock coping mechanism. The positive response suggests worsening liquidity constraints due to typhoons or excess supply of potential migrants at baseline are not binding enough to impede any average migration response from taking place.

While increasing migration rates, typhoons lead to a decrease in the wages of new migrant cohorts. Columns 2 to 5 of Table 1.2 presents the results. In the short run, a one standard deviation average hurricane exposure in the past two years decrease the average log contract

[28]Specifically, the homogeneity if treatment effect paths across treatment timing-groups corresponds to Assumption 3 in Sun and Abraham (2021a) and Assumption 6(b) in Callaway et al. (2021). Under such homogeneity (along with other assumptions of no anticipation and parallel trends), the TWFE estimates recover a weighted average of causal response parameters across units.

Table 1.2: Typhoons Increase Migration and Decrease New Migrant Wages

	(1)	(2)	(3)	(4)	(5)
		Migrant Wages ...			
	Migration rate	mean ln(wage)	ln 25th pct.	ln 50th pct.	ln 75th pct.
$T_{m,[t,t-1]}$ ($\beta_{ShortRun}$)	1.244***	-0.010***	-0.010***	-0.016***	-0.013**
	(0.413)	(0.003)	(0.003)	(0.003)	(0.005)
	[0.413]	[0.002]	[0.003]	[0.004]	[0.004]
$T_{m,[t-2,t-3]}$ ($\beta_{MediumRun}$)	1.463***	-0.005***	-0.006**	-0.010***	-0.005
	(0.492)	(0.002)	(0.003)	(0.004)	(0.003)
	[0.433]	[0.002]	[0.003]	[0.003]	[0.003]
Observations	15,970	15,788	15,768	15,768	15,768
Adjusted R2	0.877	0.904	0.676	0.837	0.863
Mean Dep. Var.	33.644	5.476	5.287	5.377	5.576

Notes: Unit of observation is municipality-year. All regressions include unit and year-by-island-group fixed effects. Migration rate is calculated per 10,000 capita. Observation numbers for columns 2-5 are lower due to municipality-years with no migration. In columns 2-5, observations are weighted by the number of migrants making up each cell. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered SEs.

wages of new migrants by 1%. The wage effect falls to about half its initial magnitude two to three years after exposure, unlike the persistent effects on migration rates.[29][30] Columns 3-5 show that the short-run drop is not constrained to the mean, with the 25th, 50th and 75th percentiles of wages falling 1%, 1.6%, and 1.3% respectively. Overall, for each one percent increase in migration in the short-run, mean (median) wages fall by 0.27% (0.43%). Again, to assess dynamics, Figure 1.2b shows event study results for average log wages.[31] The drop in migrant cohort wages is immediate, bottoms out about a year and a half after the typhoon, and dissipates faster than the migration rate results.

[29] I focus on the mean log wages throughout the body of the paper, as opposed to log of mean wages, as it yields more precise results and is consistent with individual contract level regressions I run below with $ln(wage)$ as the outcome. Column 4 of Appendix Table A.3 shows that the short-run and medium-run drop in migrant wages is around 1.5% and 0.6% with $ln(meanwage)$ as the outcome.

[30] The contract data only reports the legally contracted wage for each contract. There is anecdotal evidence of illegal practices such as recruitment agencies informing migrants that they should expect a lower wage upon arrival than the legally binding minimum that is shown in the contract (Agunias, 2010). Unfortunately, it is not possible for me to ascertain prevalence of such activity. If a natural disaster makes migrants more likely to accept such offers due to lower reservation wages, the effects obtained from the administrative data would be biased towards zero.

[31] Given the average becomes noisier when calculated over low number of observations, I calculate the wage results in bi-quarterly periods to avoid exceedingly low number of observations per cell.

Figure 1.2: Event Study Results for Migration Rate and Average Migrant Wages

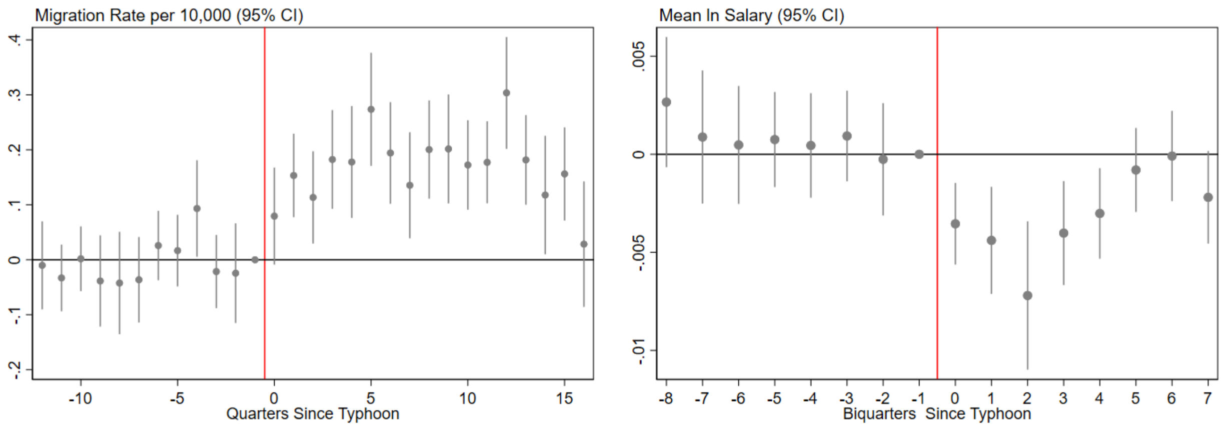

(a) Migration Rate

(b) Mean ln(Wage)

Notes: Panel (a): Dependent variable is migration rate. Unit of observation is municipality-quarter. The specification includes municipality and quarter-by-island-group FEs. Confidence intervals based on province clustered standard errors. Panel (b): Dependent variable is mean log(wage) of migrants. Unit of observation is municipality-biquarter. The specification includes municipality and biquarter-by-island-group FEs. Observations are weighted by the number of migrants in each cell. Confidence intervals based on province clustered standard errors.

Robustness. Typhoons may impact per capita migration by changing the population level or growth in affected regions, implying results might be driven by changes in the denominator as opposed to migration. Columns 1-2 of Appendix Table A.3 shows that the migration results hold with the 2000 or 2007 population as the denominator, indicating that population responses to typhoons are not driving the results. Additionally, as shown in column 3, results are broadly unchanged if I use log migrant counts as the outcome variable. A related concern is whether migrants are forced out of their origin-municipalities before migrating. The administrative data reports origin information at time of migration. Therefore some observations may have misclassified origin municipalities if they internally migrated out of a typhoon-affected municipality before migrating. This would bias my results towards zero, implying reported effects are possibly underestimates.[32]

Appendix Table A.4 further shows that the migration and wage results are robust to including year as opposed to year-by-island fixed effects or controlling for linear trends in

[32] A possible check for this concern would be replicating the analysis with municipality of birth of migrants to assess similarity of results. Unfortunately data on birth municipalities are not available.

baseline characteristics.[33] Further, given heterogeneity in average typhoon exposure and migration intensity of municipalities, I check if results are driven by an outlier region by running the main analyses while leaving out one of 79 provinces at a time. Results presented in Appendix Figure A.4 suggests no individual province is driving the results. Finally, to assess whether extreme typhoon exposure realizations are driving results, Appendix Table A.5 shows results are robust to winsorizing the exposure measure at 99%.

Spillovers. If control municipalities nevertheless are indirectly affected, my estimates would be biased due to stable unit treatment value assumption (SUTVA) violations as my research design compares compares migration outcomes across municipalities.

There are two particular concerns in this setting. First, economic and population ties between municipalities can propagate the economic impacts of typhoons to control municipalities. To assess this possibility, I use the insight from economic geography literature that population and economic ties between regions tend to fall with distance. Replicating the analysis with varying the size of the administrative units or interest (from 1597 municipalities to 17 regions) and checking if control municipalities respond to nearby typhoon exposure, I find suggestive evidence that there may be positive migration responses from nearby municipalities (Appendix Section A.5.5.1). This would imply the above migration results are under-estimates, though the estimates are imprecise.

Second, if individuals from different municipalities are competing for the same set of overseas contracts, control municipalities may see a *decrease* in migration due to increased competition.[34] I look for evidence of this by leveraging variation in which municipalities tend to send migrants to similar destinations or use the same recruitment agencies, with the idea that these municipalities would be competing for similar overseas contracts. Checking for migration response when municipalities with similar destination or recruitment agency shares are affected by typhoons, I do not find any evidence of negative spillovers. The details of the analysis is in Appendix Section A.5.5.2.

[33] Baseline municipality controls are calculated using the 2000 Census and include baseline population, share of population with primary school, secondary school, and college education, and share of households that are rural. I additionally construct province level income controls from the 2003 FIES that includes the logs of average household expenditures, average household income, and the variance of household income.

[34] Reallocating overseas contracts to regions with negative shocks would still be welfare enhancing, though would have different implications about how much aggregate migration from the Philippines have increased.

1.5.4 Migrant Wages Fall Due to Country and Occupation Downgrading

Next, I turn to the mechanisms behind the drop in migrant wages following typhoons. Migrant cohorts following typhoons have higher share of migrants leaving for lower paying occupations and destination countries. This country and occupation "downgrading" drives the drop in migrant wages. Compositional changes of migrant cohorts along observable dimensions does not explain this pattern, with typhoons *increasing* the educational attainment of migrant cohorts.

Typhoons Lead to Country and Occupation Downgrading. Figure 1.3 plots the coefficients of interest from the summary specification 1.8 with the share of migrants going to an occupation or country quartile (as discussed in Section 1.4.1) as the outcome. The left panel shows a clear pattern for destination countries: after a typhoon, share of migrants going to the lowest paying countries persistently increase while the share going to the highest paying countries decrease. A one standard deviation typhoon exposure leads to share of migrants going to the lowest wage quartile countries increasing by 1.4 and 1.8 percentage points in the short and medium run (1.9% and 2.4% of mean), while share going to the highest quartile falls by 0.9 and 1.3 percentage points (8% and 12% of mean).

A similar pattern holds for occupations in the short run. A one standard deviation typhoon exposure increases the short run share of migrants going to lowest quartile of occupation by 0.8 percentage points (1.4% of the mean), while decreasing the second quartile by 0.8 percentage points (4.5% of the mean) with no significant changes for third and fourth quartiles. However, in the medium run the pattern is flattened, with a slight increase in the share of migrants going to the higher paying occupations. The flattening of the occupation patterns in the medium run is consistent with the finding that the drop in average wages are halved in this time frame.

Downgrading Explains the Wage Decrease. To assess how much of the wage drop is driven by shifts in occupation and country shares of migrant cohort, I use the individual level contract data to check if controlling for predicted wages of a migrant based on their occupation and destination country dampens the coefficient on typhoon exposure. I estimate:

$$\ln w_{iodmt} = \alpha + \beta_{SR} T_{m(t,t-1)} + \beta_{MR} T_{m(t-2,t-3)} + \beta_{pred} \ln w_{od}^{pred} + \gamma_{r(m),t} + \gamma_m + \epsilon_{iodmt} \quad (1.10)$$

where i is an individual migrant, o and d are the occupation and destination country of the

Figure 1.3: Typhoons Increase Migration to Lowest Paying Countries and Occupations

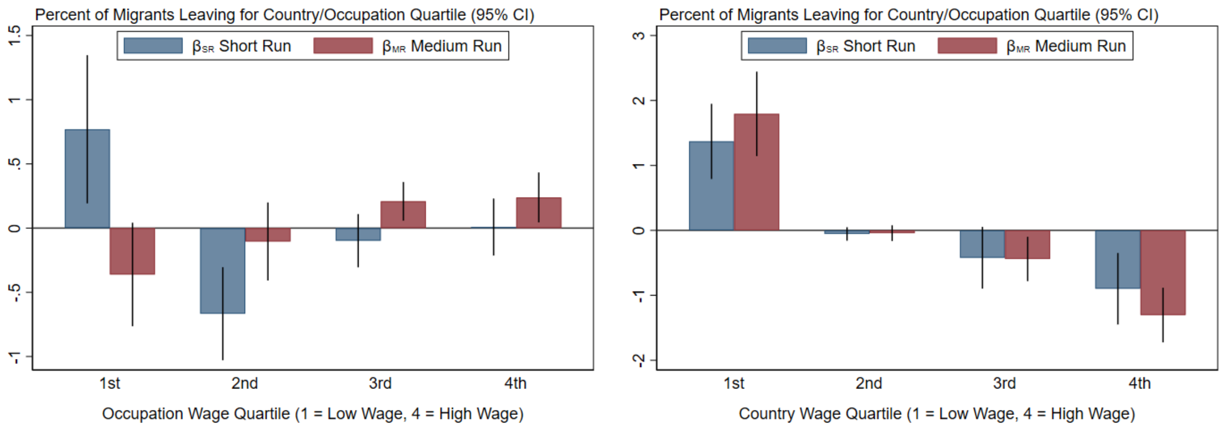

(a) Country Quartiles

(b) Occupation Quartiles

Notes: Unit of analysis is municipality-year. β_{SR} and β_{MR} from summary specification 1.8 are plotted. Outcome variable is the share of migrants leaving for overseas jobs in the specified country or occupation quartile. Estimates plotted from 8 regressions with share of migrants going to each occupation and destination quartile as the outcome. 1st quartile indicates lowest wage and 4th quartile indicates highest wage occupations and destinations. Municipality and year-by-island-group fixed effects are included. Observations are weighted by the number of migrants in each cell. Confidence intervals based on province clustered standard errors. The regression tables underlying the figure is presented in Appendix Table A.6.

migrant, m is the origin municipality, and t is year. The key addition to the municipality-level specification is to control for predicted wage based on occupation and destination $\ln w_{od}^{pred}$. To ensure predicted wages are not influenced by typhoons shocks in the analysis period, I calculate them using occupation and country fixed effects estimated on the baseline period data.[35] Results are presented in columns 1 (without predicted wage control) and 2 (with predicted wage control) of Table 1.3. The negative wage impact of typhoons vanish with predicted wage controls included, implying that the drop in wages are entirely driven by the changes in occupation and destination shares of workers.[36]

Downgrading is not Driven by Observable Selection. While the drop in migrant cohort wages can be rationalized by falling reservation wages, one potential concern is whether the effect is driven by selection. If typhoon induced migrants are negatively selected, the shift towards lower paying countries and occupations may due reflect selection as opposed

[35] Specifically, I use the sum of the estimated fixed effect for the occupation and destination of the contract (appropriately shrunk by the empirical bayes method of Morris (1983)) as a control variable. Results are almost identical if I construct predicted wages using more granular occupation-by-destination fixed effects.

[36] Columns 1-3 of Appendix Table A.7 presents analogous results controlling for occupation-by-country fixed effects as opposed to controlling for predicted wages calculated from the baseline period.

Table 1.3: Wage Effects Controlling for Occupation, Destination, and Demographics

	(1)	(2)	(3)	(4)	(5)	(6)
	Migrant Level				Municipality Level	
	Dependent Variable = ln(wage)				Log Avg. Age	Share Male
$T_{m,[t,t-1]}$ ($\beta_{ShortRun}$)	-0.010***	-0.001	-0.009***	-0.010***	-0.002***	-0.005
	(0.002)	(0.001)	(0.002)	(0.002)	(0.001)	(0.003)
					[0.001]	[0.003]
$T_{p,[t-2,t-3]}$ ($\beta_{MediumRun}$)	-0.005***	0.002*	-0.005***	-0.006***	-0.001	0.003
	(0.002)	(0.001)	(0.002)	(0.002)	(0.001)	(0.003)
					[0.001]	[0.003]
$\ln w_{od}^{pred}$		0.763***				
		(0.028)				
Demographic Controls	No	No	Yes	Yes	-	-
Demo. Cont. X Muni.	No	No	No	Yes	-	-
Observations	3,637,967	3,637,967	3,663,428	3,661,074	15,747	15,757
Adjusted R2	0.063	0.487	0.113	0.125	0.782	0.942
Mean Dep. Var.	5.527	5.527	5.530	5.530	3.458	0.361
SD Dep. Var.	0.471	0.471	0.474	0.474	0.053	0.184

Notes: Unit of observation is individual contracts for columns 1-4 and municipality-year for columns 5-6. Typhoon exposure index is at the municipality level. All regressions include municipality and year-by-island-group fixed effects. Demographic controls include sex dummies interacted with 5-year age bin dummies. Province clustered standard errors in parenthesis (79 clusters). Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered SEs.

to falling reservation wages. To assess this possibility, I focus on whether composition of migrants in terms of age, sex, and education changes after typhoon exposure.

The contract data contains information on the age and sex of the migrants. Columns 5 and 6 of Table 1.3 present results from the summary specification for the age and sex composition of migrants. While the migrant cohorts following typhoons are younger, the magnitude of the effect is small and unlikely to be economically meaningful, with a short run decrease of 0.2% for the average migrant age. There is no significant effect on the sex composition of migrants.

To assess whether this small compositional change can explain the drop in migrant wages following typhoons, I use the individual level contract data to check if controlling for age and sex of the individual migrants dampens the negative wage effects of typhoons.[37] Column 3 of Table 1.3 shows that the inclusion of these demographic controls have a negligible effect on coefficients β_{SR} and β_{MR}. In column 4, I further allow the impacts of demographic observables to vary by municipality, to allow for the possibility that, for example, younger women may be more likely to be nurses as opposed to domestic helpers in different municipalities. I again find almost identical coefficients, implying compositional changes with regards to sex and age are not a driver of the decrease in migrant wages.

Educational attainment is a key proxy for the skill and earning potential of migrants. Since the contract data lacks educational information, I turn to the Survey on Overseas Filipinos (SOF). SOF surveys a nationally representative sample of households about members who left for overseas employment in the past five years. It includes the educational attainment and year of departure of migrants. Further, from 2011 to 2016, SOF also allows identification of new land-based migrants, allowing me to focus on the same population as the main analyses. The SOF only has region level identifiers, which is a larger administrative unit that partitions the country to 17 regions. I construct a region-year level measure of educational attainment of new migrant cohorts to see if educational attainment of cohorts change in response to typhoons.[38]

Figure 1.4 presents event-study results for the share of migrants with some post-secondary education and the share completed college. Typhoons lead to an *increase* in the educational attainment of new migrants. The year following a one SD typhoon exposure, the percent of migrants who completed some secondary schooling and completed college is 3 and 2.5 percentage points higher.[39] Appendix Figure A.5 presents analogous results using the 2007, 2010, and 2015 100% census microdata on the educational attainment of the *stock* of overseas workers at the municipality level. Consistent with the SOF analysis, the educational attainment of the stock of migrants are higher in the two years following typhoons, precisely

[37] I create 5 year age bin dummies interacted with sex as controls to flexibly control for age and sex of migrants. While the variance decomposition analysis suggests that these demographic controls explain little of the variation in the presence of country and occupation fixed effects, demographics may nevertheless be important determinants of occupation and country choice, therefore meaningfully impacting the cohort wage distribution.

[38] The region level typhoon exposure index is created analogously to the municipality-level index, except aggregated from pixel-level to region-level. Note that the the coarseness of the identifiable geography decreases power and likely biases the estimates towards zero due to the measurement error.

[39] The less pronounced effects on the same year as typhoon exposure likely partially reflects that the SOF is conducted in October, before the completion of the typhoon season.

when we see the most pronounced drops in new migrant wages.[40]

Figure 1.4: Typhoon Driven Migrants are More Educated

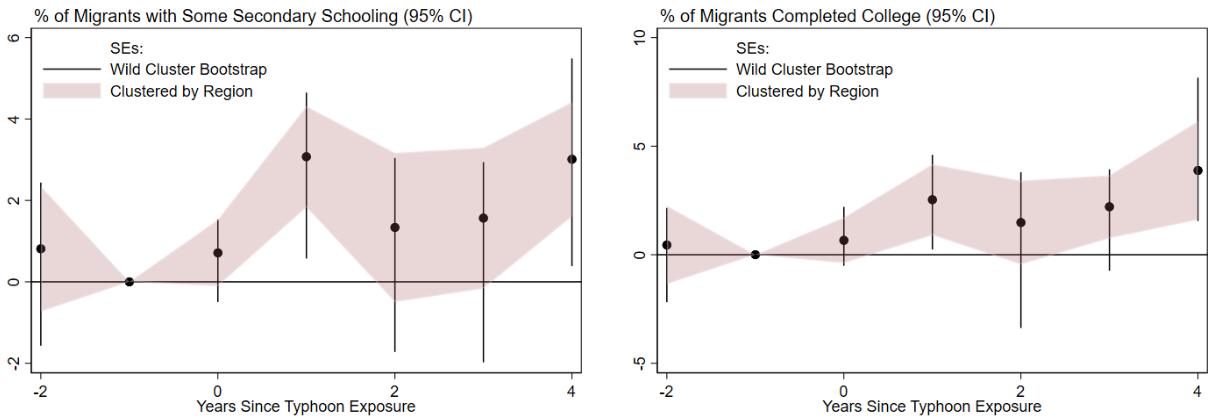

(a) New Migrants: Share with Post Secondary (b) New Migrants: Share Completed College

Notes: Unit of analysis is region-year. Outcome variables are the educational attainment of the new migrants, constructed from the 2011-2016 Survey on Overseas Filipinos. Region and year-by-island-group fixed effects are included. Confidence intervals calculated by wild-cluster bootstrap due to low number of clusters ($N_{cluster} = 17$). Observations are weighted by number of migrants making up each cell.

This finding is consistent with the well established competing forces of high returns to migration versus liquidity constraints in the context of international migration from developing countries (Mckenzie and Rapoport, 2007; Bazzi, 2017). If typhoons increase the return to migration, but also makes the liquidity constraints more binding for lower wealth households, we would expect to see that share of new migrants with higher wealth or income would increase. As wealth is associated with educational attainment, the marginal migrant in response to a typhoon would have higher educational attainment than the average. Further, given there exists competition for overseas contracts, an increase in highly educated *potential* migrants may lead to disproportionately more highly educated actual migrants if they out-compete their lower educated counterparts for the overseas opportunities. Both these possibilities imply that, while access to overseas work allows for some shock-coping in aggregate, it can also contribute to uneven recovery as individuals with higher education can take advantage of these opportunities.

Overall, it is unlikely that negative selection explains the drop in wages in response to

[40]Note that the educational attainment of the stock of labor migrants would not only be impacted by the educational attainment of new land-based migrants, but also by the return decisions of migrants and educational attainment of sea-based migrants. Given data limitations, I can't However, two estimates from different sources providing qualitatively similar estimates strongly suggest typhoons increase the educational attainment of new migrants.

typhoons. If anything, migrant cohorts following typhoons have higher educational attainment, which should, all else equal, lead to higher wages in the migrant cohort.[41] Of course, in the absence of individual panel data, this analysis cannot rule out unobservable differences that may make typhoon induced migrants have a preference or ability towards lower paying occupations and destination countries.

1.5.5 Alternative Explanation for Wage Drop

The characteristics of the migration response are precisely what one would expect if congestion and search frictions induce migrants to lower their reservation wages and possibly direct their search effort towards lower paying overseas markets with higher likelihood of securing a job. I consider the plausibility of two potential alternative mechanisms below.

Equilibrium Wage Changes in Destinations. Can the estimated wage effects be explained by a drop in equilibrium destination wages due to increased migration from the Philippines? Given the change in occupation and destination shares explain the estimated wage drop (Table 1.3 columns 1-2), this can't be the driver of the observed wage effects. Note that, to drive the negative wage estimates, any destination wage response have to be municipality specific as more aggregate wage responses would be subsumed by the year-by-island-group fixed effects. This implies that the wage drops estimated from my micro research design would underestimate the total wage effects if typhoons lead to an unlikely aggregate change in Philippine migrant wages.[42]

Increased Demand for Local Construction Workers. In response to a typhoon, the share of migrants going to the occupations in the second wage quartile experience the most pronounced drop. This group of occupations include many potentially construction related occupations such as "concrete finishers", "carpenters", and "roofers". This raises the question of whether the occupation downgrading is driven less by changes in reservation wages and more due to an increased demand for local construction work in response to destruc-

[41] While the SoF does not provide information on migrant earnings, it does provide information about remittances sent by migrants and which regions the migrants are in. College educated land-based migrants, on average, send 42% more cash remittances, are 2.3 times (3.8 pp) more likely to be in Europe or the Americas (destinations with high earnings), and 25% (8 pp) less likely to be in Gulf countries (destinations with low earnings) compared to migrants with only high school completion.

[42] Wage changes in destinations are unlikely given migration from affected municipalities constitute a small fraction of the labor force in destination countries. Appendix Table A.2 shows the share of Filipino migrants in the migrant stock of the top-20 destination countries for 2015. The share is never over 13% (Brunei) and is approximately 4% for the median top destination country. Considering the share of Filipinos in the total labor force is smaller than their share in the total migrant stock, it is unlikely that an out-migration increase of approximately 4% from an affected region will affect destination wages meaningfully. Further, migrant wages are highly regulated, with much of migration taking place in overseas labor markets with binding wage regulation (see Section 1.2.1 and McKenzie et al. (2014)).

tion caused by typhoons. In Appendix Table A.7 column 4, I test directly whether share of migrants leaving for construction related occupations fall in response to a typhoon, and find small and insignificant effects. In columns 5-6, I check if controlling for the share of migrants leaving for construction occupations dampen the drop in migrants leaving for the second wage quartile. The coefficient falls by about 20% (which is expected as about 60% of migrants in the second quartile are in construction related occupations), but is still substantial and significant. Therefore, the occupational downgrading does not seem to be driven by increased demand for (potentially skilled) construction workers in origin municipalities.

1.5.6 Interpreting Effect Sizes

The evidence presented so far shows that typhoons lead to increased out-migration and a drop in average average migrant cohort wages. Taking the estimated coefficients as causal estimates and focusing on the entire typhoon exposure in the data, I find that typhoons lead to about 10,100 additional migrants each year to leave the Philippines. This corresponds to a 2.7% increase in aggregate migration from the Philippines each year. The combined contract earnings of typhoon induced migrants are around 129,184,710$ (ignoring the possibility that subset of the contracts will likely be renewed in completion), which corresponds to about 20% of the annual typhoon damage estimates by the government in this period.[43] Finally, if the average wages of migrants did not fall in response to typhoons, the total earnings would have been 22% higher at 157,678,510$, ignoring the possibility that the higher educated marginal migrants could have secured higher paying contracts in the absence of downgrading. This underlines that potential lost income due to downgrading can be significant. Research that only focus on migration rates without taking migrant wages into account can significantly overestimate the aggregate migrant earning response to origin shocks. The details of these calculations can be found at Appendix Section A.4.1.

1.6 Effects of Destination Country Demand During Typhoon Shocks

The ability to respond to shocks through migration is dependent on availability of overseas jobs. The abundance of overseas contracts when a typhoon hits can therefore mediate the

[43]Damage estimates take into account damage to building, infrastructure, and the estimated value of agricultural loss. While 20% is a substantial share of damages, note that only a subset of the total migrant earnings are saved or remitted back to the Philippines, and the domestic productivity of migrants are now not realized because they are work abroad.

migration response. In this section, I examine the heterogeneity of migration response to typhoons by the international migrant demand conditions facing a municipality.

1.6.1 Measurement of Migrant Demand Conditions

To assess the effects of migrant demand conditions facing a Philippine municipality in a year, I create a time variant municipality-level migrant demand proxy. The proxy follows a shift-share structure. It combines information about baseline migrant shares to each destination country (shares) with plausibly-exogenous labor demand conditions at destination countries proxied by real GDP per capita (shifts). Note that my interest is not in the direct effects of the proxy itself. I instead focus on the interaction between the typhoon shocks and the migrant demand proxy to analyze differential responses to typhoons due to concurrent migrant labor demand conditions.

I first define $\pi^0_{m\to d}$ as the share of baseline period migrants from municipality m going to destination d. As shown in Appendix Figure A.3, there is substantial variation across municipalities in which destination countries migrants tend to go to. I then use the (lagged) real GDP per capita of destination countries $\ln GDP_{m,t-1}$ as a proxy for the migrant demand conditions at each country. With these ingredients, the province level proxy is defined as the migrant share weighted average GDP of destination countries:

$$D_{m,t} = \sum_d \pi^0_{m\to d} \, \ln GDP_{d,t-1} \tag{1.11}$$

As I discuss further below, the mean difference in $D_{m,t}$ across municipalities, driven by baseline migration shares $\pi^0_{m\to d}$, are not assumed to be exogenous. Therefore, I demean the migrant demand index within municipalities to ensure that the variation driving estimation is the year-to-year variation. Two main requirements for the *relevance* of this proxy for migration response is that (1) baseline destination shares are persistent enough to be operational in the analysis period and (2) destination GDP per capita is a meaningful measure of migrant demand conditions in each destination. Failing these requirements would bias coefficients towards zero. I assess these requirements next.

Migrant Shares are Persistent. Appendix Figure A.6 demonstrates that baseline migrant shares $\pi^0_{m\to d}$ are highly correlated with the analysis period migrant shares $\pi^{2007-2016}_{m\to d}$, with a Pearson's correlation coefficient of 0.81. Demeaned by destination country, the correlation is still 0.57. Such persistence is consistent with prior evidence showing migration flows are channelled between local areas and destination countries, due to network facilitation of migration, information frictions, and specialization of recruitment agencies for specific destination markets over time (Cortes, 2015; Munshi, 2003a; Khanna et al., 2022; Shrestha and Yang, 2019b).[44]

Migration to a Destination is Increasing in Destination GDP. Appendix Section A.5.4 shows that total migration from the Philippines to a destination is increasing in the GDP per capita of the destination. I find that the destination GDP elasticity of migrant flows from the Philippines is ranging from 1.5 to 4 across specifications. However, there is no significant effect on migrant wages. The increase in migration without an increase in wages is consistent with excess migrant supply in prevailing overseas wages, where positive destination GDP shocks increase the quantity of overseas jobs available for Philippine municipalities and slackens the migrant labor market (McKenzie et al., 2014).

In Appendix Section A.5.4.2, I further assess whether migrant flows *in response to typhoon shocks* are differential in destination country GDP. Using a bilateral province-destination country specification, I find that the migration response to typhoons from a province to a destination country is increasing in destination country GDP per capita. While suggestive, this shouldn't be interpreted as direct evidence of a larger regional migration response based on demand conditions. With a bilateral specification, we cannot rule out the possibility that migrants from a municipality may just be reallocating across destinations based on relative abundance of overseas contracts across countries. For the effects of migrant demand conditions on total municipality-level migration response, I turn to the main analysis below.

1.6.2 Empirical Approach

To examine the impact of contemporaneous migrant demand conditions on the short-run migration responses to typhoons, I estimate the following specification:

$$y_{mt} = \beta_1 T_{m,(t,t-1)} \times D_{m,t} + \beta_2 D_{m,t} + \beta_3 T_{m,(t,t-1)} + \delta X_{mt} + \gamma_m + \gamma_{r(m),t} + \epsilon_{mt} \tag{1.12}$$

where y_{mt} is the outcome of interest for municipality m in year t. $D_{m,t}$ is the (demeaned) migrant demand index. As before, $T_{m,(t,t-1)}$ is the average typhoon exposure T of municipality m in the past two years, γ_m is municipality fixed effects that controls for any time-invariant characteristics of the geographic unit of analysis, and $\gamma_{r(m),t}$ is an island-group by year fixed effect that controls flexibly for differential trends in the outcome within Philippine island

[44]For example, the 2015 and 2016 KNOMAD/ILO surveys suggest about how half the migrants learned about their current overseas job from relatives or friends. Similarly, in a 2004 survey, 67% of first time migrants report knowing a member of their social network in their destination (Cortes, 2015).

groups over time. X_{mt} includes additional controls, including past typhoon exposure. The coefficient of interest is β_1, which captures whether responses to typhoons are differential in contemporaneous migrant demand conditions.

It is critical to ensure that β_1 is not driven by trends in $D_{m,t}$ common to all municipalities due to global shocks to or trends in GDP. For example, because GDP of all destination countries tend to be higher by the end of the analysis period, β_1 may erroneously pick up differing migration responses to typhoons at the end of the analysis period for reasons unrelated to demand conditions. As discussed in the inference section below, I generate counterfactual municipality-year level demand indices by permuting the residual growth rate of countries controlling for global GDP shocks and possible persistence of growth rates over time. The expected value of the counterfactual demand indices, $\mathbf{E}[\tilde{D}_{mt}]$, capture the expected evolution of $D_{m,t}$ taking global trends, shocks, and persistence of growth rates into account. I control for the direct effect and the typhoon exposure interaction of $\mathbf{E}[\tilde{D}_{mt}]$ in the empirical specification. Controlling for $\mathbf{E}[\tilde{D}_{mt}]$ ensures that the coefficient of interest is identified from deviations from the expected value of $D_{m,t}$ across municipalities, as opposed to the overall time trend of $D_{m,t}$ common to all municipalities.[45] I additionally show results further controlling for (1) the interaction between typhoon exposure and a linear time trend, (2) interactions between typhoon exposure and global GDP per capita, Philippines GDP per capita, and a weighted average of all destination countries GDP per capita (weighted by national baseline migration shares), and finally, most stringently, (3) interaction of typhoon exposure with year dummies.

Identification. The identifying variation for the interaction term stems from the year to year variation in D_{mt}. The key identifying assumption is that the year-to-year variation in destination country GDP per capita, i.e. the yearly GDP growth, is as-good-as-random from the perspective of individual Philippine municipality migration decisions (Adao et al., 2019a; Borusyak et al., 2022b). This exogeneity condition is sound as destination countries' GDP per capita are driven by a myriad of factors independent of Philippine municipality-level migration such as global economic conditions, local consumer spending and business decisions, and international oil prices for gulf countries. With the Philippine migration from typhoon affected municipalities constituting a vanishing fraction of destination labor supply and my focus on new contracts, any concerns about reverse causality due to typhoons are

[45] As discussed below, my setup is analogous to the non-random exposure (baseline migration shares) to exogenous shocks (destination GDP changes) setting described in Borusyak and Hull (2020). In such settings, controlling for the mean counterfactual shock is essential to avoid omitted variable bias stemming from the non-random exposure to shocks. Controlling for $\mathbf{E}[\tilde{D}_{mt}]$ serves essentially the same purpose in my setting. Furthermore, permuting residual growth rates allows for randomization inference that avoids potentially misleading conventional standard errors due to possibly correlated residuals between municipalities with similar exposure shares. I further describe the procedure below.

highly diminished. To further ensure that any reverse causality between regional Philippine migration and contemporaneous economic activity in destination countries is not driving results, I lag the destination demand index by one year.

Exclusion Restriction. My goal is to isolate the effects of migrant demand conditions a municipality faces on the migration responses to typhoons. The relevant "excludability" criteria is that the baseline migration shares are not correlated with other economic ties between destination countries and Philippine municipalities. For example, if the export or the FDI share of a municipality-country pair are highly correlated with migration shares, $D_{m,t}$ could conceivably affect migration outcomes through income effects (Orefice et al., 2023). Due to a lack of municipality-level trade and FDI data, I am unable to check for the correlation between the shares directly.[46] Instead, I directly test whether the migrant demand index is associated with domestic income using the Family Income and Expenditure Surveys. Appendix Table 1.6 shows results from a household-level panel regression of domestic (non-remittance) income on the province level demand index D_{pt}.[47] There is no significant relationship between domestic income or any of its sub-components and D_{pt}. This alleviates concerns that the contemporaneous migrant demand influences domestic income of Philippine municipalities directly.

Inference. Conventional standard errors in shift-share designs can be misleading due to possibly correlated residuals between municipalities with similar exposure shares (Adao et al., 2019a; Borusyak et al., 2022b). Given that my empirical setting can be construed as a case of non-random exposure to exogenous shocks, I turn to the randomization inference procedure of Borusyak and Hull (2020).[48]

The procedure requires generating counterfactual migrant demand indices D_{mt}. Given

[46] Appendix Table A.11 reports aggregate statistics on the top export destinations and top sources of FDI for the Philippines for the analysis period. Top 20 export and FDI destinations, accounting for 94 % and 95% of exports and FDI flows, only accounts for 16% and 25% of temporary labor migration out of the Philippines in this period. Therefore, the scope for positive income gains through export market expansion or increased FDI seems limited from GDP growth in migration destination countries. For imports, the share is up to 82%, which is driven by Gulf countries being both common destination countries and large exporters of oil. The differential exposure of Philippine municipalities to oil imports would primarily be driven by the industry structure of the municipality, as opposed to the share of their temporary migrants going to gulf countries.

[47] The household survey data does not allow the analysis to be done at the municipality level. I construct the province level migrant demand index analogously to the municipality level.

[48] The exposure-robust standard error procedure of Borusyak et al. (2022b) is not suited to my setting for two reasons. First, I am interested in the coefficient estimated on the interaction between the demand index measure and the typhoon shock, which is not supported by the method. Second, while there are numerous destination countries for Philippine migrants, the high concentration of migration in a relatively low number of destination countries violates a key "law of large numbers" assumption for the Borusyak et al. (2022b) procedure to be asymptotically valid. The randomization inference procedure introduced by Borusyak and Hull (2020) remains valid in the presence of concentrated exposure.

the yearly GDP per capita change of a country is assumed to be as-good-as-random from the perspective of municipality level migration outcomes, I first create counterfactual annual GDP growth for each destination country. I estimate the following AR1 specification for log GDP per capita growth rates for country c in year t on the panel of destination countries:

$$\ln(g_{ct}) = \theta \ln(g_{ct-1}) + \delta_t + \epsilon_{ct}$$

where δ_t captures the annual averages due to global shocks and the AR1 term captures any potential persistence in growth rates. I generate counterfactual growth rates by varying the error term to ensure global shocks and structural persistence are preserved. Specifically, I re-sample year-country level innovations $\hat{\epsilon}_{ct}$ from the empirical distribution of $\hat{\epsilon}_{ct}$ within each year and generate counterfactual GDP growth $\tilde{g}_{ct}$ using the randomized error terms and model estimates.[49] Using the counterfactual growth rates, I generate country-level counterfactual GDP per capita using observed 2000 GDP per capita as the baseline. Combining the counterfactual destination GDP with baseline migrant shares generates the counterfactual demand index $\tilde{D}_{mt}$.

I run the estimating equation 1.12 with the counterfactual $\tilde{D}_{mt}$ and check if the estimated $\tilde{\beta}_1$ and $\tilde{\beta}_2$ have larger magnitudes than $\hat{\beta}_1$ and $\hat{\beta}_2$. I repeat this procedure 1000 times and display the fraction of regressions where counterfactual estimates are larger in magnitude than main estimates. I present these p-values alongside province-clustered and Conley standard errors.

1.6.3 Results

Table 1.4 presents the results, with all the right hand side variables are normalized to have zero mean and unit standard deviation. Column 2 reports results from the preferred specification with $\mathbf{E}[\tilde{D}_{mt}]$ controls. As expected, better demand conditions are associated with more migration on average (row 2). More importantly for the purposes of this paper, there is a clear relationship between migration response to typhoons and migrant demand in destination countries: better demand conditions lead to substantially larger migration response without a correspondingly higher drop in migrant wages. A standard deviation increase in the demand index from the mean approximately doubles the short-run migration response to

[49] In Appendix Table A.12, I show robustness of the main results to modifying the inference procedure by (1) including an AR2 term, (2) including country fixed effects, (3) including an interaction between year and gulf country dummies to capture possibly correlated shocks across these common migrant destinations, and (4) combining all three. I further relax the assumption that the error terms are distributed from the empirical distribution and show robustness to instead assuming they are normally distributed with mean and variance matching the distribution of empirical error terms for each year in the analysis.

typhoon exposure (panel A, column 2), while dampening the wage drop by about 20% (column 2 of Panel B, imprecisely estimated). Appendix Table A.12 shows that the interaction results are robust to alternative randomization inference procedures.

Additionally, Columns 3 and 4 shows that the migration response is robust to the inclusion of linear time trend and additional global and domestic GDP controls. Column 5 reports results from the most stringent specification with typhoon exposure interacted with year dummies. The coefficient is still positive and large for migration rate, and is highly statistically significant with the randomization inference procedure, but the conventional standard errors are larger due to reduced variation.[50]

Figure 1.5a visualizes the migration rate results by using the estimated model in column 2 of Table 1.4 to trace out the average short run migration response to a one standard deviation typhoon exposure across 10th to 90th percentile of migrant demand index. The migration response (as percent of the mean migration rate) increases steeply with the migrant demand index. The response is small and statistically insignificant below the 35th percentile, is 3.5% at the median, and climbs to 6.7% at the 80th percentile.

Next I turn to the average migrant cohort wage response to typhoons per additional (percent) migration. Panel B of Figure 1.5 presents the results.[51] Overall, better migrant demand conditions decrease the typhoon induced migration elasticity of migrant wages. In other words, better demand conditions not only increase the migration response, they also dampen the per migrant drop in new contract wages. The average drop in migrant wages per percent migration falls from 26% in median demand conditions to 11% at the 80th percentile.

The dampened migrant wage response is primarily driven by lower occupational downgrading during better demand conditions. Figure 1.6 presents analogous figures for the share of migrants leaving for each occupation quartile (Appendix Table A.13 shows the regression results underlying the figure). The jump in the share of migrants going to lowest paying occupations is decreasing in the migrant demand index, while the fall in the share going to higher paying occupations are increasing, with the effect most pronounced for the second quartile, which exhibits the biggest fall after a typhoon. When typhoons coincide with better migrant demand, typhoon induced migrants (who are on average more educated) are better able to secure higher paying occupations, decreasing the pressure to occupationally downgrade. In contrast to the occupation results, Appendix Figure A.7 shows that country downgrading in response to typhoons does not fall during better demand conditions. This is due to the fact

[50] Note that the coefficient on the typhoon exposure cannot be interpreted as the average effect of one standard deviation typhoon shock when typhoon $\times$ year dummies are included as controls.

[51] Results are shown for 40th to 90th percentile of the migrant demand index, because the small estimated migration response below this point leads to extreme and highly imprecise estimates as the migration response is at the denominator.

Table 1.4: Migration Response to Typhoons are Larger During Better Demand Conditions

	(1)	(2)	(3)	(4)	(5)
Panel A. Outcome: Migration Rate (per 10,000)					
$T_{m,[t,t-1]}$	1.144**	1.101**	1.112***	1.012***	10.641**
	(0.434)	(0.429)	(0.364)	(0.344)	(4.354)
	[0.475]	[0.474]	[0.436]	[0.403]	[4.055]
$D_{m,t}$	2.245***	5.058***	5.060***	5.166***	5.154***
	(0.674)	(1.279)	(1.280)	(1.265)	(1.316)
	[0.650]	[1.196]	[1.195]	[1.180]	[1.205]
		{0.019}	{0.018}	{0.014}	{0.015}
$T_{m,[t,t-1]} \times D_{m,t}$	1.244**	1.233**	1.276**	1.810**	1.744
	(0.591)	(0.588)	(0.569)	(0.749)	(1.275)
	[0.678]	[0.669]	[0.693]	[0.789]	[1.143]
		{0.066}	{0.015}	{0.008}	{0.013}
Observations	13,550	13,550	13,550	13,550	13,550
Mean Dep. Var.	36.822	36.822	36.822	36.822	36.822
Panel B. Outcome: ln(Mean Wage)					
$T_{m,[t,t-1]}$	-0.009***	-0.010***	-0.010***	-0.010***	-0.003
	(0.002)	(0.002)	(0.003)	(0.002)	(0.057)
	[0.002]	[0.002]	[0.003]	[0.003]	[0.052]
$D_{m,t}$	-0.008*	-0.010	-0.010	-0.009	-0.009
	(0.004)	(0.007)	(0.007)	(0.007)	(0.007)
	[0.004]	[0.008]	[0.008]	[0.008]	[0.008]
		{0.624}	{0.588}	{0.623}	{0.622}
$T_{m,[t,t-1]} \times D_{m,t}$	0.004*	0.002	0.001	-0.009	-0.005
	(0.002)	(0.002)	(0.004)	(0.006)	(0.012)
	[0.002]	[0.002]	[0.004]	[0.006]	[0.011]
		{0.716}	{0.877}	{0.173}	{0.524}
Observations	13,516	13,516	13,516	13,516	13,516
Mean Dep. Var.	5.485	5.485	5.485	5.485	5.485
$\mathbf{E}[\tilde{D}_{mt}]$ control	No	Yes	Yes	Yes	Yes
Typh X Linear Trend	No	No	Yes	No	No
Typh X Controls	No	No	No	Yes	No
Typh X Year FE	No	No	No	No	Yes

Notes: Unit of observation is municipality-year. All regressions include municipality fixed effects, year-by-island-group fixed effects, and a control for past typhoon expsure. Municipalities with no migration in baseline period are dropped from the analysis due to missing baseline shares. Observation numbers for Panel B are lower due to municipality-years with no migration. Panel B observations are weighted by the number of migrants making up each cell. Columns 2-5 include the direct effect and the interaction with typhoon exposure of mean counterfactual demand indices. Column 3 includes a control for typhoon exposure interacted with linear time trend. Column 4 includes controls typhoon exposure interacted with Philippines GDP, global GDP, and a weighted average of all destination country GDPs (weighted by baseline national migration shares to each country). Column 5 includes typhoon exposure interacted with year fixed effects. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. Randomization inference p-values in curly brackets where applicable. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered SEs.

Figure 1.5: Migration and Migrant Wage Responses Along Migrant Demand Index

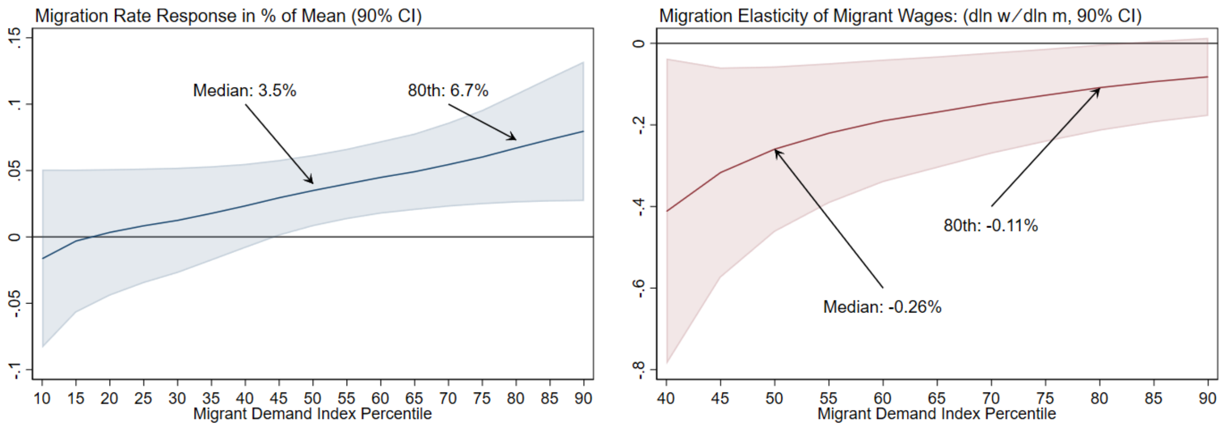

(a) Migration Response　　(b) Migration Elasticity of Migrant Wages

Notes: The unit of analysis is municipality-year. Estimation uses demeaned migrant demand index with mean counterfactual demand index controls, corresponding to columns 3 and 6 of Table 1.4. For panel (a) traces out the migration rate response divided by mean migration rate. Standard errors are clustered at the province level. Panel (b) traces out log mean wage response divided by the percent migration response to typhoons. Migration responses small for migrant demand index, therefore, estimates and standard errors of panel (b) gets extreme for low values of migrant demand index (small denominator) and are not shown. Province-clustered standard errors are shown where standard errors in (b) are calculated using the delta method.

that the migrant demand index is constructed using all destination countries, both high- and low-wage. If better demand conditions are driven by GDP growth in low-wage countries, migration response to shocks could still be biased towards these countries.

Taken together, the ability of Philippine municipalities to use international labor migration as an shock coping mechanism is strongly mediated by the international migrant demand conditions they face. Better demand conditions, through increasing overseas job availability, leads to a bigger migration response with a lower relative drop in migrant wages. Concretely, total migrant earnings response to a one standard deviation typhoon shock is about 6 percentage points in the 80th percentile demand conditions as opposed to 2.6 percentage points in the median, a substantial difference of a factor of 2.3.[52]

Placebo Exercise. To ensure that I am not merely capturing unobserved trends in migration responses to typhoons, I replicate the analysis with future and past values of the migrant

[52] Migrant earnings are given by average migrant wages $\bar{w}$ times the number of migrants m. Therefore, the total (new) migrant earning response to typhoons is given by $(1 + \frac{d\ln \bar{w}}{d\ln m})d\ln m$. During median demand conditions we get a $(1 - 0.26) \times 0.035 = 2.6\%$ response, while at 80th percentile demand conditions a $(1 - 0.11) \times 0.067 = 6\%$ response.

Figure 1.6: Occupation Share Response to Typhoons Along the Migrant Demand Index

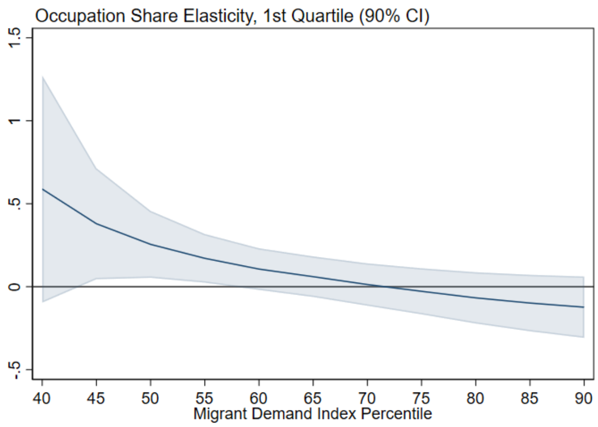

(a) 1st Occupation Quartile

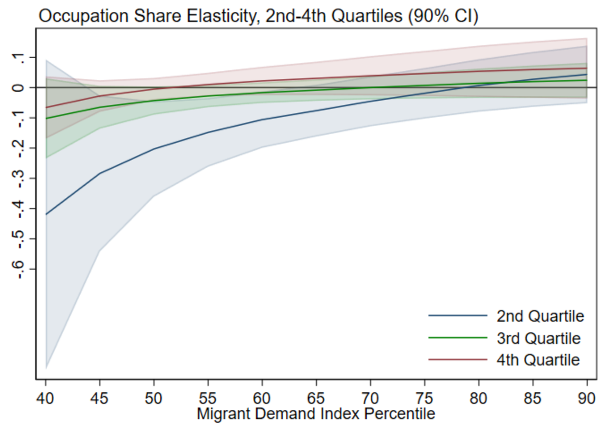

(b) 2nd-4th Occupation Quartiles

Notes: The unit of analysis is municipality-year. Estimation includes the mean counterfactual demand index controls, corresponding to specification in columns 2 and 4 of Table 1.4. Regression results underlying the figures can be found in Appendix Table A.13. Each panel traces out the occupation share response heterogeneity for share of migrants going to a different quartile. Standard errors are calculated using the delta method using the variance-covariance robust at clustering at province level.

demand index on the right hand side. Appendix Figure A.8a plots the coefficients on the interaction term of interest for migrant demand index, with migration rate as the outcome. Reassuringly, future and past levels of network GDP has a much weaker and insignificant relationship with migration responses to typhoons. Because GDP levels are persistent across time, I also undertake a placebo exercise with the migrant demand index calculated using GDP growth as opposed to levels. Appendix Figure A.8b shows that destination GDP growth in the past two years predict a stronger migration response to typhoons, while future GDP growth does not.[53]

1.6.4 Discussion and Policy Implications

This section provides among the first empirical evidence that the ability to utilize international labor migration as an ex-post coping mechanism is tempered by global migrant demand conditions the origins face at the time. Given the level of economic activity in destination countries is not a policy lever for sending countries, what are the implications of

[53]The analysis period includes the great recession. While this is not an identification concern, if the results are driven purely by the variation induced by the great recession, the external validity of the results in stable global economic conditions may be suspect. However, results are broadly unchanged if year 2008 and 2009 are dropped from the analysis.

these finding for policy?

Broadly, the evidence suggests contemporaneous increase in the availability of migrant contracts following a negative shock can lead to a stronger migration response. Therefore, policies that increase the availability of overseas jobs to affected communities in the wake of shocks can have significant shock-coping benefits. Such policies can take many forms, including the relevant government agencies intensifying efforts to secure overseas contracts, providing incentives for private recruitment agencies to have more contracts available, reducing the costs of migrating (for example through subsidizing recruitment or processing fees), and increasing access to already available overseas jobs through job fairs. In the wake of the catastrophic 2013 Typhoon Haiyan, there are reports of the Filipino government and recruiting agencies following such policies, including the organization of job-fairs in affected areas that include overseas jobs, the government securing additional overseas contracts, and recruitment agencies waiving recruitment fees. Results are also suggestive that there may be gains from diversifying the "portfolio" of foreign markets available to Filipinos regions, as this would decrease the dependence on any individual destination country.

These findings are also suggestive of how policies that restrict migration in destination countries can create significant negative externalities in terms of shock-coping for origin countries with strong migration ties. From the perspective of a Philippine province, a drop in availability of foreign contracts would have similar consequences whether it is due to a drop in a prominent destination's economic activity, or because more restrictive policies are enacted in the destination. In the wake of such restrictions, the capacity of a Philippine municipality to cope with a disaster could be greatly diminished, at least in the short run before additional adjustments can take place. This is particularly relevant given the political debates that are calling for restrictive barriers to migration (Cinque and Reiners, 2023).

It is harder to draw firm conclusions about the impacts of a more *permanent* increase in availability of overseas contracts from the empirical results at hand, as the analysis is focused on responses to changes migrant demand in the short-run. The effects therefore may partially be driven by increased search activity in the wake of typhoons leading to improved learning about the migrant demand conditions. In the absence of a shock, such shifts in migrant demand could have translated to migration outcomes slower due to information frictions (Porcher, 2022). In the case of permanent increases to availability of contracts and a long enough time horizon that such information is fully internalized, there would be a shift in the baseline rate of migrants in the economy. This can potentially lead to a more robust response from migrants that are already abroad at the time of a shock, though whether it would also lead to stronger migration response is unclear. Of course, persistent increase in access to overseas occupations can have effects on the origin economy beyond just migration levels in the long run, as documented in, for example, Khanna et al. (2022). Insofar as such increases lead to an increase in wealth available at origin over time, these regions can better invest in mitigation and adaptation technologies beyond international migration.

1.7 Remittence Response to Typhoons

International migration aids origin shock-coping primarily through remittances sent home. Much of temporary international migration is particularly motivated by supporting family members and social networks back home, and temporary migrants tend to save and remit more of their foreign earnings (Dustmann and Mestres, 2010; Yang, 2011).[54] Therefore, I conclude the paper by studying the remittance responses to typhoons using household survey data. Beyond documenting the average remittance response to typhoons, my goal is to assess the importance of new migration for the remittance response by (1) showing that the remittance response is larger during good migrant demand conditions, which should primarily act through increasing new migration, and (2) providing a back of the envelope calculation of how much of the remittance response can be attributed to the new migration response.

1.7.1 Remittance Response to Typhoons

For the remittance analysis, I use the province level typhoon exposure measure (which is calculated analogously to the municipality level) as the Family Income and Expenditure Surveys (FIES) does not allow for consistently identifying the municipalities of households and is not representative at the municipality level. I estimate the following specification at the household level, analogous to the summary specification 1.8:

$$rem_{hpt} = \alpha + \beta_1 T_{p(t,t-1)} + \beta_2 T_{p(t-2,t-3)} + \delta' \mathbf{x_{hpt}} + \gamma_{rt} + \gamma_p + \epsilon_{ipt} \tag{1.13}$$

where the outcome is household level for remittance outcome rem_{hpt} of household h, in province p, in year t. A possible concern is that the composition of households included in the FIES may respond to typhoon exposure. While I do not find imbalances in likely time-invariant household level covariates in response to typhoons (Appendix Table A.8), I nevertheless include a vector household controls $\mathbf{x_{hpt}}$ which includes household size and the

[54] According to 2006-2018 FIES 95% of Philippine households with a member currently working overseas receive income from abroad. The 2015/2016 KNOMAD/ILO surveys find that the average Philippine migrant in the sample remit around 67% of their foreign earnings.

demographics of the household head such as sex, age, and age squared, and the education level of the household head. Households are weighted by the provided sampling weights.

Table 1.5 presents the results, with odd columns excluding and even columns including household level controls. Typhoons increase per capita household remittances. Focusing on column 2, a one standard deviation exposure increases per capita remittance receipts by 384 PhPs in the short run, corresponding to 6.7% of the mean. Column 4 shows that the results are similar when using the cubic root of remittance per capita as the outcome variable to adjust for the right skew of remittance per capita data.[55] Column 6 shows a one standard deviation typhoon exposure increases the share of households receiving any remittances by 0.8 percentage points (3% of the mean). The medium run remittance results are positive as well, with magnitudes 65 to 80% of short run results when household controls are included, yet they are imprecisely estimated and statistically insignificant at 5%.

Table 1.5: Typhoons Increase Remittances

	(1)	(2)	(3)	(4)	(5)	(6)
	Abroad Inc Per Cap.		(Abroad Inc Per Cap.)$^{\frac{1}{3}}$		**1**[Any Abroad Inc.]	
$T_{p,[t,t-1]}$ ($\beta_{ShortRun}$)	329.686**	380.646***	0.247***	0.290***	0.006	0.008**
	(127.538)	(109.670)	(0.092)	(0.084)	(0.004)	(0.003)
	[154.598]	[140.258]	[0.133]	[0.126]	[0.005]	[0.005]
$T_{p,[t-2,t-3]}$ ($\beta_{MediumRun}$)	237.281	252.732*	0.188	0.203	0.007	0.007
	(164.306)	(140.932)	(0.145)	(0.128)	(0.006)	(0.006)
	[170.597]	[147.142]	[0.146]	[0.130]	[0.006]	[0.005]
HH Controls	No	Yes	No	Yes	No	Yes
Observations	306,315	306,315	306,315	306,315	306,315	306,315
Clusters	79	79	79	79	79	79
Mean Dep. Var.	5646	5646	6.3	6.3	0.27	0.27
SD Dep. Var.	15340.428	15340.428	11.285	11.285	0.443	0.443

Notes: Household level regression using 2006, 2009, 2012, 2015, and 2018 FIES data. Unit of observation is a household. Typhoon exposure is at province-year level. All regressions include province and year-by-island-group fixed effects. Observations are weighted by the provided sampling weights. Household controls are household size, gender of HH head, age (and age squared) of HH head, and whether the HH head completed primary school, secondary school, some college, or college. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered SEs.

Heterogeneity by Demand Conditions. To assess whether the stronger migration in times of high migrant demand translates into a stronger remittance response, I add the interaction between the typhoon exposure with province-level migrant demand conditions.

[55]Due to high incidence of 0s (74% of households in the sample report no remittances), I do not log-transform remittances and use the cubic root transformation instead.

Table 1.6 shows that the remittance response patterns that are parallel to that of migration response. During better migrant demand conditions, remittance response to typhoons are stronger. With the counterfactual demand index controls included, a one standard deviation improvement in demand leads to 67% increase in the short term remittance response ($\frac{276.72}{415.48}$, column 2), and almost doubles the increase in the likelihood a household receives any remittances (column 6).

Table 1.6: Remittance Response Heterogeneity by Migrant Demand Index

	(1)	(2)	(3)	(4)	(5)	(6)
	Abroad Inc Per Cap.		(Abroad Inc Per Cap.)$^{\frac{1}{3}}$		$\mathbf{1}$[Any Abroad Inc.]	
$T_{p,[t,t-1]}$	339.25***	415.48***	0.28***	0.34***	0.010**	0.011*
	(83.88)	(113.44)	(0.08)	(0.11)	(0.004)	(0.005)
	[91.99]	[107.20]	[0.09]	[0.11]	[0.004]	[0.005]
$\ln D_{p,t}$	536.90***	562.13***	0.37**	0.40**	0.008	0.009
	(197.88)	(202.16)	(0.18)	(0.18)	(0.008)	(0.008)
	[158.85]	[163.16]	[0.15]	[0.15]	[0.007]	[0.007]
	{0.424}	{0.284}	{0.520}	{0.388}	{0.684}	{0.556}
$T_{p,[t,t-1]} \times \ln D_{p,t}$	385.91***	276.72**	0.33***	0.24*	0.012***	0.009
	(98.13)	(123.91)	(0.09)	(0.12)	(0.004)	(0.006)
	[94.83]	[116.67]	[0.08]	[0.11]	[0.004]	[0.005]
	{0.020}	{0.104}	{0.024}	{0.086}	{0.024}	{0.078}
$\mathbf{E}[\tilde{D}_{mt}]$ control	No	Yes	No	Yes	No	Yes
Observations	306,315	306,315	306,315	306,315	306,315	306,315
Mean Dep. Var.	5646	5646	6.3	6.3	0.27	0.27

Notes: Household level regression using 2006, 2009, 2012, 2015, and 2018 FIES data. Unit of observation is a household. Typhoon exposure is at province-year level. All regressions include province and year-by-island-group fixed effects and household controls. Household controls are included: household size, gender of HH head, age (and age squared) of HH head, and whether the HH head completed primary school, secondary school, some college, or college. Observations are weighted by the provided sampling weights. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. Randomization inference p-values in curly brackets. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered SEs.

1.7.2 Share of Remittance Response Due to New Migration

The aggregate importance of the migration response is partly related to how much of the increase in remittances can be attributed to new migration versus increased remittances from already existing migrant networks. The FIES does not provide information on the source of remittances, rendering it impossible to directly assess how much of the remittance response can be attributed to new migrants. I proceed with two back of the envelope calculations to assess what share of the remittance response can be reasonably attributed to new migration, with each calculation following a different strategy employing different assumptions. Appendix Section A.4.2 provides the details of these back of the envelope calculations.

Both strategies require comparable estimates of the migration response and the remittance response to typhoons. Because the remittance responses can't be estimated at the municipality level and administrative data does not allow for household level analysis, Appendix Table A.9 shows the province-year level remittance per capita, migration rate, and average migrant wage response estimates underlying the back of the envelope calculations.[56]

Strategy 1: Comparing levels of remittance and migrant earnings response. I first directly compare the level of remittance response to the migrant earning response, using the estimates from regressions reported in Appendix Table A.9. Between 2007-2016 the short-run and medium-run total remittance responses to typhoons are 458,737,800$ and 390,660,700$ in an average year. It is well established that remittance levels as reported by households can be underestimates, with De Arcangelis et al. (2023) calculating that households under-report remittances by 23% in the Philippines. Therefore, I scale up the remittance responses by 23% to 595,763,300$ and 507,351,600$. The total annual migrant earning response in an average year is 68,482,200$ and 60,065,450$ in the short- and medium-run. Because vast majority of migrants leave for 2 year contracts, I multiply these estimates by two as typhoon-induced migrants from the previous year would also continue to remit. Assuming 67% of migrant earnings are remitted each year[57], I estimate 15% of short-run and 16% of medium-run response is attributable to new migration taking place due to typhoons. Panel A of Table 1.7 summarizes the calculations.

Strategy 2: Comparing percent remittance and migrant earning response. Let total annual remittances $R = r \times y_M$ be the remittance rate r times the total migrant earnings subject to being remitted y_M. To a first order, percent change in total remittances due to a typhoon is:

$$\underbrace{d\ln R}_{\substack{\text{Total Remittance}\\ \text{Response}}} = \underbrace{d\ln r}_{\substack{\text{Remittance Rate}\\ \text{Response}}} + \underbrace{share^{new}}_{\substack{\text{New Migration Share}\\ \text{in Total } y_M}} \times \underbrace{d\ln y_M^{new}}_{\substack{\text{Migrant Earning}\\ \text{Response}}}$$

where y_M^{new} is the total earnings of *new migrants* and $share^{new}$ is the share of new migrant

[56]Because the household level estimates with provided survey weights weighs each province by its population size, the analysis presented in Appendix Table A.9 furthers uses population weights for provinces.

[57]The value is calculated from the 2015 and 2016 KNOMAD/ILO Migrant Cost Surveys which surveys 849 Filipino labor migrants in Qatar and UAE across variety of occupations. The average migrant reports remitting 67% of their annual income home.

earnings in total migrant earnings. I calculate the percent change in new migrant earnings in response to typhoons directly from the migration response estimates ($d \ln y_M^{new}$) in Appendix Table A.9. To proxy $share^{new}$, I need two pieces of information. First, what share of remittances are from temporary labor migrants, as opposed to larger diaspora including permanent migrants out of the Philippines? Using aggregate data on the sources of total cash remittances to the Philippines in 2013 from the central bank, I estimate around 41% of remittance flows are from temporary labor migrants.[58] Second, what is the share of *new* temporary labor migrants among the stock in a given year? Using census and survey data on the stock of migrants, I estimate this share as 40.5%.[59] Adjusting for the subset of new migrants who are seafarers as opposed to land-based migrants (the focus of my analysis), I conclude $share^{new} = 0.108$. With these values, 16% of short-run and 18% of medium-run remittance results are plausibly attributable to new migration. Panel B of Table 1.7 summarizes the calculations.

Table 1.7: Share of Remittance Response Attributable to New Migration Response

	$d\,\$R$	$r \times d\,\$y_M^{new}$	$d \ln R$	$share^{new}$	$d \ln y_M^{new}$	Share
Strategy 1 - Levels						
Short-Run	595,763,300\$	45,883,074\$	-	-	-	15%
Medium-Run	507,351,600\$	40,243,851\$	-	-	-	16%
Strategy 2 - Percent Change						
Short-Run	-	-	7.2%	0.108	5.8%	17%
Medium-Run	-	-	6.0%	0.108	5.3%	19%

Notes: Appendix Table A.9 columns 1-2 presents the province-level estimates that underlie $d \ln y_M^{new}$ and column 3 presents the province-level estimates that underlie $d \ln R$ (remittance response). New migrant earning response $d \ln y_M^{new}$ is calculated as the province-level percent migration response times the average new migrant wage response to typhoons. Share indicates the share of total migration response attributable to new migrant response, calculated as $2 \times \frac{r \times d\,\$y_M^{new}}{d\,\$R}$ for Strategy 1 and $2 \times \frac{share^{new} \times d \ln y_M^{new}}{d \ln R}$ for Strategy 2. Multiplication by two captures that migrants predominantly leave for two year contracts. New migrants from the previous year would still be sending remittances in the year of interest. All underlying estimates are weighted by population province.

Overall, both calculations suggest a substantial share of the remittance response can be

[58]Specifically, I combine information from the central bank about the source regions of remittances to Philippines with 2013 Commission on Filipinos Overseas estimates of what share of migrants in each region are temporary labor migrants as opposed to permanent migrants. Details are in Appendix Section A.4.2.

[59]This number is consistent with the fact that 94% of overseas contracts in the data have a duration of less then three years, with two year contracts by far the most common. Therefore it is to be expected that a large share of total migrants earnings is from relatively new contracts/migrants. To calculate the proxy, I calculate migrant stock estimates from census microdata and the government Survey on Overseas Filipinos (SOF) reports, and use yearly migrant flow data from POEA reports. Using census microdata for 2007, 2010, and 2015 I find an average new migrant share of 44%. Using 2007-2016 SOF reports, the new migrant share is 37%. I use the average of these two estimates.

attributable to new migration taking place in response to typhoons. The migration response is therefore not only critical for the households that now have access to remittances, but also is important for the aggregate origin province level shock-coping. Note that three assumptions implicit in these calculations suggest that these figures are likely underestimates for the contribution of the migration response to the aggregate remittance response: (1) Typhoon-driven migrants have the same remittance rates as the average migrant, (2) return decisions of previous migrants or seafarer migration are not affected by origin typhoon exposure, and (3) none of the short-run migration episodes induced by typhoons are affecting the medium-run remittances (i.e. all migration episodes are 2 years or less).

1.8 Conclusion

I investigate how temporary international labor migration responds to origin shocks, and how this response is mediated by the demand conditions facing potential migrants. Focusing on a decade of typhoon shocks in the Philippines and using administrative data on new migrant contracts, I find that a 1 standard deviation typhoon exposure increases out-migration from a region by about 4% for up to 3 years. However, migrant cohorts following typhoons have lower wages primarily due to going to lower paying occupations and countries, even though such cohorts have higher educational attainment. This suggests that search frictions in these markets lead to occupation and country downgrading in response to negative home shocks. Better contemporaneous demand conditions when shocks occur strengthen the migration response significantly, and dampens the wage drop. Remittance responses to typhoons mirror these patterns, with about 67% larger remittance due to a 1 standard deviation increase migrant demand conditions at the mean.

Given the policy debates surrounding how to facilitate and regulate this increasingly common form of migration, these findings have important policy implications. Results suggest access to such international labor markets can have shock-coping benefits, even in the presence of well documented frictions. Such benefits should be considered in assessing the benefits and costs of promoting international labor migration. Further, destination policies that restrict such migration can impose negative externalities on origin countries with strong ties by decreasing the origin countries' ability to dampen the impacts of negative shocks. These findings gain additional relevance given the extreme weather events are expected to increase over time due to climate change.

CHAPTER 2

Abundance from Abroad: Migrant Income and Long-Run Economic Development

with Gaurav Khanna, Caroline Theoharides, and Dean Yang

2.0 Abstract

How does income from international migrant labor affect the long-run development of migrant-origin areas? We leverage the 1997 Asian Financial Crisis to identify exogenous and persistent changes in international migrant income across regions of the Philippines, derived from spatial variation in exposure to exchange rate shocks. The initial shock to migrant income is magnified in the long run, leading to substantial increases in income in the *domestic* economy in migrant-origin areas; increases in population education; better-educated migrants; and increased migration in high-skilled jobs. 77% of long-run income gains are actually from domestic (rather than international migrant) income. We empirically demonstrate that these findings are not confounded by potential trade impacts of the same exchange rate shocks. A simple model yields insights on mechanisms and magnitudes, in particular, that 23.2% of long-run income gains are due to increased educational investments in origin areas. Improved income prospects from international labor migration not only benefit migrants themselves, but also foster long-run economic development in migrant-origin areas.

2.1 Introduction

Moving from a developing to a developed country for work leads to income gains that are larger than the impacts of any known economic development program (Clemens et al., 2019; Pritchett and Hani, 2020). International migrants from developing countries sent home $669 billion in remittances in 2023, an amount as large as all foreign direct investment, and more than three times larger than foreign aid flows to the developing world (World Bank, 2023).[1] Motivated by these economic gains, most developing-country governments have policies facilitating international migrant labor (United Nations, 2019b).

There is ample evidence that international migration raises incomes for the migrants themselves. However, evidence is scarce on how international migrant income affects broader economic development in migrant-origin *areas*. Positive shocks to the income of international migrants could loosen liquidity constraints on human capital and entrepreneurial investments in origin areas. In addition, higher *potential* income in the international labor market could have effects even in households initially without migrants, by raising the returns to migration. As a result, migration rates could rise. Furthermore, households could invest more in education, because education raises the likelihood of securing an overseas job, and also has returns in overseas work. Increases in such investments in migrant-origin areas should raise longer-run economic growth. Evidence of such development impacts would suggest that international migration policies could play a more prominent role in efforts to reduce global poverty (Nunn, 2019).

We ask how persistent increases in international migrant income affect long-run economic development in migrant-origin areas. We exploit a large-scale natural experiment: persistent changes in international migrant incomes across Philippine migrant-origin areas driven by persistent exchange rate changes due to the 1997 Asian Financial Crisis. Philippine provinces varied prior to 1997 in the amount of migrant income earned by their citizens in many different countries. The vast majority of these migrant workers were overseas on temporary labor contracts (returning eventually to their origin areas). Overseas migrant income sources then experienced exogenous – and heterogeneous – exchange rate shocks in 1997, which persisted. To undertake our analyses, we obtained unusual Philippine government administrative data on migrant worker contracts, with information on migrant incomes, origin provinces, and overseas destinations. The combination of the natural experiment and these unique data makes possible a shift-share identification strategy. We examine aggregate impacts on 74 Philippine provinces up to two decades later.

Our empirical analyses implement frontier methods for identification and inference in shift-share research designs, following Borusyak et al. (2022a). Each province's exposure "shares" are pre-shock levels of migrant income per capita from each international migrant destination (which we call "exposure weights"). These exposure weights vary greatly across origin provinces and overseas destinations. For example, 1995 migrant income emanating

[1] International migration also involves large numbers of people. 210 million people from developing countries were international migrants in 2019 (United Nations, 2019a), a magnitude similar to the number of microcredit clients, 140 million (Convergences, 2019), or conditional cash transfer (CCT) program beneficiaries, 185 million (World Bank, 2018b).

from Japan is 10.7 times higher on a per capita basis for Bulacan province (PhP 3,540 per provincial resident) than for Leyte (PhP 332 per provincial resident).[2] Japan's exchange rate shock should therefore have 10.7 times greater impact on population-level mean outcomes in Bulacan than in Leyte.

Each destination's "shift" is its exchange rate shock. Table 2.1 displays the exchange rate shock for the top 20 migrant destinations in the immediate post-shock year (1997-1998). These exchange rate movements were persistent over the next two decades, as we discuss further in Section 2.4.4. The shocks range from a 4% depreciation against the Philippine peso for Korea to a 57% appreciation for Libya. Other important destinations such as Japan and Taiwan fall in between (32% and 26% appreciations, respectively). The identification assumption is that these exchange rate shocks are as-good-as-randomly assigned. Balance tests with respect to pre-shock characteristics support this identification assumption.

We present the resulting variation in the shift-share variable across provinces in Figure 2.1. The shift-share variable is interpreted as a shock to migrant income per capita (i.e., per provincial resident). We estimate the impacts of this shock on long-run provincial outcomes. Impacts could be due to the positive income shock experienced by migrants who were overseas when the shock occurred. Households initially without migrants at the time of the shock could also change their migration decisions and education investments in response to the increase in the return to migration. Standard errors account for correlation of shocks across provinces with similar exposure weights (Borusyak et al., 2022a).

We find, first, that the initial shock to migrant income (measured by our shift-share variable) is magnified over time. Each unit short-run (1997-1998) positive shock to migrant income is increased more than five-fold in the longer run (through 2009-2015). Below, we explore the mechanisms behind this substantial magnification in the context of a structural model.

Second, we find that the positive migrant income shocks lead to substantial increases in *domestic* Philippine income per capita (not including migrant income or remittances) in migrants' origin provinces. A province's "global income" per capita is the sum of its domestic income and (international) migrant income per capita. 77% of the long-run increase in global income per capita is from the increase in domestic income, and 23% is from migrant income. We also see corresponding increases in household expenditure per capita. These gains emerge over roughly two decades after the 1997 shocks, reflecting persistence in the exchange rate changes and in the overseas sources of migrant income for particular Philippine provinces. The magnitude of the gains is nontrivial. A one-standard-deviation shock raises

[2] All Philippine peso (PhP) amounts in this paper are in real 2010 pesos (PPP exchange rate 17.8 PhP/USD).

Figure 2.1: Spatial Distribution of Shift-Share Variable (Migrant Income Shock) Across Philippine Provinces

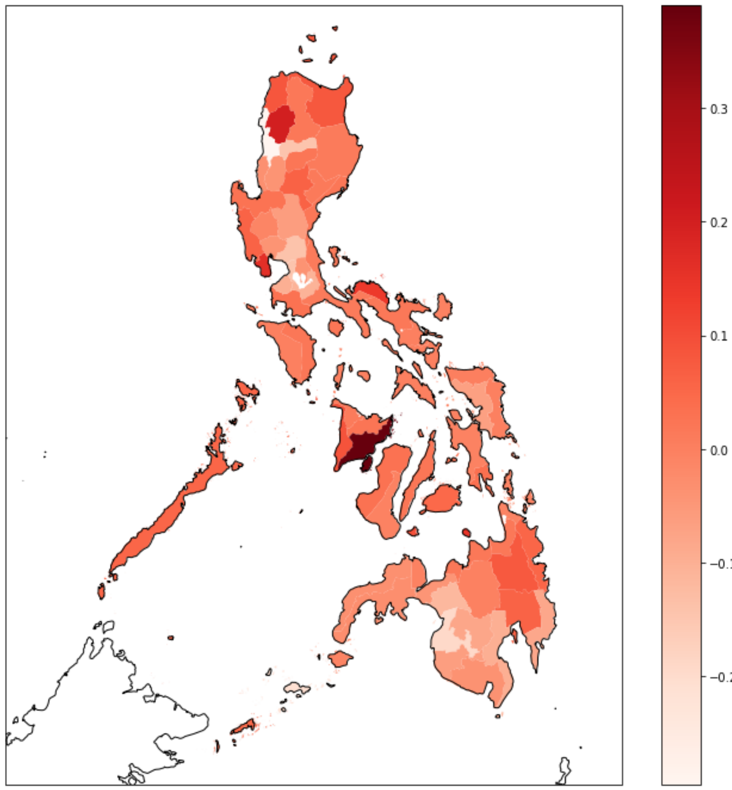

Notes: Spatial variation in province-o shift-share variable (migrant income shock) $Shiftshare_o = MigInc_{o0}Rshock_o$ after partialling-out weighted average exchange rate shock $Rshock_o$ and pre-shock migrant income per capita $MigInc_{o0}$, for 74 Philippine provinces. See Section 2.4 and Appendix Section B.2.2 for details.

global income per capita 12-18 years later by 2,272 Philippine pesos (PhP) (0.18 standard deviation).

We take seriously a range of threats to causal identification. Most prominently, it is crucial to test whether the shift-share variable on which we focus is operating via effects on international trade flows rather than (solely) effects on migrant income. Exchange rate shocks generated by the Asian Financial Crisis – the fundamental shocks driving our shift-share variable – can also clearly affect imports and exports, which in turn could also affect development outcomes in Philippine provinces. We investigate whether impacts of our migrant-income shift-share variable operate (at least in part) via impacts on international trade. We construct additional shift-share variables capturing the impact of exchange rate shocks on provincial imports and exports, in the same spirit as our migrant-income shift-share variable. The import and export shift-share variables exploit (pre-1997) variation in exchange rate shocks in import and export partners, in combination with province-level employment shares in import and export industries. Our results are highly robust to controlling for these import and export shift-share variables, which suggests that our estimates primarily reflect impacts due to changes in potential migrant income. We also provide additional evidence that province-level exports are not responsive to the migrant income shocks, and foreign direct investment (FDI) is also unlikely to be a relevant mechanism. This further helps confirm that the shift-share variable of interest operates as a shock to migrant income, rather than trade or FDI.

Throughout our analyses, we also provide two additional categories of tests of the credibility of our causal claims. First, we test whether changes in outcomes in the pre-shock period ("pre-trends") are correlated with the future value of the shift-share variable. We find no evidence of pre-trends, ameliorating concerns that provinces that would have higher values of the shift-share variable (after the 1997 Asian Financial Crisis) were already experiencing more positive trends in development outcomes even before 1997. Second, we consider potential omitted variables at the origin-province or migrant-destination level. Our estimates are generally not sensitive to controls accounting for ongoing trends or heterogeneity in exposure to the Asian Financial Crisis-induced downturn related to baseline province characteristics such as industrial structure and development status.

We provide further insights into mechanisms and effect magnitudes with the help of a simple structural model. We use the model to derive our estimating equation, quantify the contribution of various channels, and see if our framework can rationalize the magnification of the income gains. We augment a gravity model of migration (Llull, 2018; Bryan and Morten, 2019; Lagakos et al., 2023) to allow workers to make educational investments and enter skilled occupations. Persistent positive migrant income shocks may alleviate constraints on such investments, and increase the return to migration.

Given the central role of skill in the model, we empirically estimate impacts on educational investments. We find large positive effects: a one-standard-deviation migrant income shock increases the share of the population with a college education by 0.51 percentage points (0.11 standard deviation). We also show that these increases in skill in the population are accompanied by increases in the share of migrants who are college-educated, and in new labor migration in highly-skilled occupations overseas.

We estimate that educational investments account for 23% of the increase in global income per capita. Furthermore, the model fully explains the over-five-fold magnification of the effect of the shift-share shock on migrant income, derived from increases in educational investments in the population, increasing migrant skill levels, and changes in migration patterns across destinations.

We also provide a stylized framework to understand the plausibility of our estimated effects on domestic income. We make assumptions regarding the share of migrant income returned to origin economies, the aggregate demand multiplier, and the return on entrepreneurial investments. A reasonable set of such assumptions yields the observed long-run increase in domestic income.

Our study is made possible by two unusual elements. First, the natural experiment of the 1997 Asian Financial Crisis generates the exogenous exchange rate variation central to our shift-share identification strategy.[3] Second, we obtained unusual Philippine government administrative data on migrant worker contracts. Without these data, provincial exposure weights ("shares" in the shift-share) would have been unobservable, making the shift-share strategy impossible.

This paper contributes to research on the economic impacts of international migration on developing-country populations. Prior research has established causal impacts of migrant economic conditions or migration opportunities on migrants' origin *households*.[4] Our work is related to a small body of recent research on economic impacts of international migration on migrant-origin *areas*, that emphasizes causal identification. Theoharides (2020) finds that closing a prior migration opportunity reduces income and raises child labor in Philippine origin areas. Dinkelman and Mariotti (2016) and Dinkelman et al. (2024) examine long-run impacts of migrant work in South Africa on Malawian origin-area education and development. Caballero et al. (2023) study short-run effects of migrant exposure to Great Recession shocks on Mexican-origin areas.[5]

An important feature of our paper is our focus on the impacts of increased international income from formal, legal migrant labor. Unlike undocumented and unregulated migrant flows across borders, migration that is facilitated and regulated by governments is highly policy-relevant, and most developing country governments are taking concrete steps towards promoting it (as we discuss in Section 2.2). Credible evidence on the impacts of legal, regulated international migrant labor flows on origin-area economic development is of interest to development policy-makers.

This paper has several additional distinguishing features, compared to prior research.

[3]Prior studies have exploited international migrants' exchange rate shocks to study impacts on migrants and their origin households (Yang, 2006, 2008b; Kirdar, 2009; Nekoei, 2013; Abarcar, 2019; Dustmann et al., 2023).

[4]Such prior works include Dustmann and Kirchkamp (2002), Yang (2008c), Gibson et al. (2010), Gibson et al. (2011), Mendola (2012), Gibson et al. (2014), Clemens and Tiongson (2017), Gröger (2019), Cuadros-Menaca and Gaduh (2020), Mobarak et al. (2023b), and Bossavie et al. (2021b).

[5]In studies of *internal* (within-country) migration, Kinnan et al. (2019) examine impacts of Chinese migration on origin areas using an instrument based on shocks in domestic migrant destinations, and Akram et al. (2017) examine Bangladeshi village-level impacts of randomly inducing rural-urban migration.

First, we examine long-run impacts, up to two decades after the initial shock. Dinkelman and Mariotti (2016) and Dinkelman et al. (forthcoming) also estimate long-run effects. Those studies differ in estimating long-run impacts of a brief historical episode of migrant work that did not persist. We study a shock to migrant income with long-run persistence, and a migrant flow that also persists. This allows us to examine how resulting investments in education initiate a virtuous migration cycle, by enabling high-skilled future migration, with subsequent increases in future migrant income.

By exploiting persistent exogenous variation in migrant income opportunities, we are able to answer a fundamental question in the economics of migration: do origin areas with greater access to high-income migration opportunities develop faster than origin areas with less attractive migration opportunities? We are able to plausibly identify the causal impact of persistently higher migrant income opportunities, and thus reveal whether migration policy can be used effectively as a part of economic development policy.

In addition, our work is distinct in simultaneously examining impacts on migrant, domestic, and global income, due to our novel data on migrant income. We can, therefore, examine the relative magnitudes of impacts on domestic income and migrant income, and thus conclude that the vast majority of long-run gains are from increases in domestic income. Finally, we complement our reduced-form estimates with a structural approach to provide insights on mechanisms and the long-run magnification of income gains.

Our findings are reminiscent of the recent literature finding positive long-run impacts of asset transfers to catalyze income gains from household entrepreneurial enterprises (de Mel et al., 2008; Banerjee et al., 2015; Bandiera et al., 2017; Banerjee et al., 2021), and providing evidence of poverty traps (Balboni et al., 2021; Kaboski et al., 2022). The migrant income shocks we study could have long-run impacts, in part, by enabling escapes from poverty traps. Our finding that a substantial share of gains in domestic income come from household enterprises suggests that migration policy can be an effective tool in the development anti-poverty toolkit.

This paper also contributes to research on the impacts of migration on skill composition at origin. Our conclusions concord with prior findings that migration leads to "brain gain," stimulating educational investments, and raising general skill levels back home (Stark et al., 1997; Mountford, 1997).[6] These findings contrast with studies finding that migration leads to a net loss of skilled individuals from the population ("brain drain"), in part via reductions in schooling investments (McKenzie and Rapoport, 2011; de Brauw and Giles, 2017; Tang

[6] Such studies include Batista et al. (2012), Docquier and Rapoport (2012), Clemens and Tiongson (2017), Shrestha (2017), Theoharides (2018b), Chand and Clemens (2019), Khanna and Morales (2023), and Abarcar and Theoharides (2022).

et al., 2022).[7] We add to this literature by finding that increases in education may generate a virtuous cycle, leading to higher-skilled future migration, which in turn raises incomes and education levels.

2.2 Context: International Labor Migration

210 million individuals from developing countries were international migrants in 2019. The largest source countries of international labor migrants are India, Mexico, and China; Bangladesh, Pakistan, the Philippines, and Indonesia also send substantial numbers abroad (United Nations, 2019a). Moving from a developing to developed country for work is associated with substantial income gains for migrants (Clemens et al., 2019). Gibson et al. (2018), Mobarak et al. (2023b), and Gaikwad et al. (2024) find that random assignment to international migrant work opportunities leads to improved migrant income, and better outcomes for migrants and their origin households.[8] Income gains from increased international migration flows are orders of magnitude larger than the likely impacts of further liberalization of international trade or capital flows, or of *in situ* efforts to raise incomes in the domestic economy of developing countries (Clemens, 2011b; Pritchett and Hani, 2020).

Motivated by these gains, most developing country governments facilitate their citizens' international labor migration. We tabulated data on government policies on outbound international migration collected by United Nations (2019b). Out of the 70 developing countries with populations exceeding 1 million, 94% have a dedicated government agency implementing migration policy; 88% have a dedicated government agency for overseas employment, citizens abroad, or diaspora engagement; and 78% have policies promoting migrant remittances.

In the Philippines, two government agencies facilitate international labor migration. The Philippine Overseas Employment Administration (POEA) regulates international migrant recruitment, issuing operating licenses to recruitment agencies and reviewing and approving migrant work contracts. The Overseas Workers Welfare Administration (OWWA) works to ensure the well-being of overseas Filipino workers (OFWs) and their families. It intercedes (via Philippine consulates worldwide) for workers experiencing abuse or contract violations, repatriates workers in conflict zones, assists OFW families in hardship, and facilitates the return and "reintegration" of OFWs to the Philippines. POEA and OWWA are the sources

[7] Evidence on reductions in education investment due to factory openings in Mexico (Atkin, 2016) is also relevant.

[8] Moreover, many prior studies have established positive correlations between international migration and economic development outcomes in origin areas (e.g., Lopez-Cordoba (2005), Acosta et al. (2008), Orrenius et al. (2010)).

Table 2.1: Exposure Weights and Exchange Rate Shocks in Top 20 Destinations of Filipino Migrants

Destination	Mean Exposure Weight	Std. Dev. Exposure Weight	10th Percentile Exposure Weight	90th Percentile Exposure Weight	Exchange Rate Shock (1997-1998, $\tilde{\Delta}R_d$)	Exchange Rate Change, 1994 - 1996 (pre-shock)
Japan	792.10	1130.49	81.69	2326.40	0.32	–0.07
Taiwan	709.79	804.84	63.41	1872.03	0.26	–0.04
Saudi Arabia	670.42	583.41	196.61	1635.78	0.52	–0.01
Hong Kong	576.08	787.50	37.90	1640.57	0.52	–0.01
United States	452.86	509.16	48.32	1045.28	0.52	–0.01
United Arab Emirates	126.23	132.14	21.35	236.41	0.52	–0.01
Malaysia	74.56	85.63	5.30	172.55	–0.01	0.04
Kuwait	72.27	218.87	0.00	77.34	0.50	–0.02
Qatar	66.98	91.55	0.74	142.48	0.52	–0.01
South Korea	54.51	108.20	0.00	103.49	–0.04	–0.01
Brunei Darussalam	50.87	43.54	8.47	108.42	0.30	0.08
Oman	47.40	319.45	0.00	21.25	0.52	–0.01
Libya	40.85	38.73	2.64	83.48	0.57	–0.21
Guam	38.10	90.22	0.00	89.82	0.52	–0.01
Italy	30.43	55.54	0.00	100.28	0.38	0.04
Canada	29.91	44.13	0.00	84.75	0.42	–0.01
Northern Mariana Islands	28.17	40.10	0.00	73.16	0.52	–0.01
Bahrain	25.67	43.89	0.00	49.30	0.52	–0.01
Singapore	25.18	24.68	0.00	72.84	0.29	0.08
Israel	17.12	94.28	0.00	16.59	0.38	–0.06

Notes: Table displays 20 destinations d with the highest mean exposure weight (across provinces o). Columns 1-4 present summary statistics for exposure weights ω_{do0}, across 74 Philippine provinces o ("shares" of the shift-share variable). See Subsection B.2.2 and Section 2.4 for details on exposure weight definition. Columns 5 and 6 present exchange rate changes. Column 5 displays exchange rate shock $\tilde{\Delta}R_d$ ("shift" of the shift-share variable). Exchange rate shock is change in Philippine pesos (PhP) per foreign currency unit. Exchange Rate Shock (1997-1998, $\tilde{\Delta}R_d$) is fractional change between July 1996-July 1997 and October 1997- September 1998 (e.g., 10% appreciation of the foreign currency against the Philippine peso is 0.1). Column 6 (Exchange rate change 1994-1996) is corresponding fractional change in exchange rate between 1996 and 1994, before July 1997 Asian Financial Crisis. 84 additional destinations not shown.

of the migrant contract data we use in our analyses.[9]

In recent decades, increasing shares of the Philippine population have migrated, had a household member migrate, or had overseas income. From 1990 to 2015, the fraction of the

[9]There are several prominent examples of government agencies facilitating migration in other developing countries. In Pakistan, the Bureau of Emigration and Overseas Employment regulates and licenses recruitment agencies. The Ministry of Labor, Migration, and Employment of the Population in Tajikistan regulates migration and facilitates job matching. Agencies in Bangladesh (the Bureau of Manpower, Employment, and Training and the Welfare Fund for Migrant Workers) and in Indonesia (the National Authority for the Placement and Protection of Indonesian Overseas Workers) play similar roles to the Philippines' migration agencies.

population currently overseas rose from 0.7% to 2.2%, and the fraction of households with an overseas migrant member rose from 3.2% to 7.5%. The share of households with overseas income rose from 16.6% in 1991 to 29.7% in 2018.[10] The vast majority of migration outflows from the Philippines is migration for temporary, legal work by workers who expect to return to their origin areas after one or more labor contracts.

Migrant income in the Philippines comes from numerous overseas destinations, and migrant destinations vary substantially across origin provinces. Table 2.1 shows the top 20 migrant destinations, ranked by mean "exposure weight" across provinces (1995 migrant income per capita, for province-destination dyads). Our empirical approach exploits the fact that, for each destination, there is substantial variation in the exposure weight across provinces.

2.3 Data and Measurement

We summarize data sources here; details are in Appendix B.1. We examine outcomes of 74 Philippine provinces,[11] typically over triennial periods or periods determined by census rounds.

2.3.1 Construction of Shift-Share Variable

To obtain causal estimates, we exploit the component of changes in provincial migrant income per capita that is due to the 1997 Asian Financial Crisis exchange rate shocks. The shift-share variable that isolates this exogenous variation in provincial migrant income per capita is our causal variable of interest.

$Shiftshare_o$ is the predicted short-run change in migrant income per capita due to the exchange rate shocks. In Appendix B.2.2 we derive this shift-share variable from a simple theoretical model of migration, which we then use to quantify mechanisms and gauge plausibility of effect magnitudes.

The "exposure weight" ω_{do0} serves as the "share" in the shift-share. ω_{do0} captures the extent to which a typical province-o resident is exposed to a destination-d exchange rate shock. ω_{do0} is province o's pre-shock aggregate migrant income from destination d, divided by province population to yield a per capita measure.

[10] Overseas income is primarily migrant remittances, but also includes sources such as pensions and investment income.

[11] To deal with changes in provincial definitions and borders, we combine geographic areas and work with a consistent definition of 74 provinces with borders as they were defined in 1990.

The "shifts" in the shift-share are the destination-d exchange rate shocks $\tilde{\Delta} R_d$. Exchange rate shocks $\tilde{\Delta} R_d$ affect a province-o resident in proportion to the magnitude of migrant income per capita coming from destination d prior to the crisis; we thus refer to the ω_{do0} terms as "exposure weights".[12]

To calculate province o's shift-share measure, each destination-d exchange rate shock $\tilde{\Delta} R_d$ is multiplied by the corresponding exposure weight ω_{do0}, and then summed across destinations d. $Shiftshare_o$ is thus the predicted change in province-o migrant income per capita due to the exchange rate shocks:

$$Shiftshare_o = \sum_d \left(\omega_{do0} \tilde{\Delta} R_d \right) \tag{2.1}$$

Now, multiply and divide $Shiftshare_o$ by the pre-shock sum of migrant income across destinations ($\sum_d \omega_{do0}$, the sum of exposure weights). This yields the following expression, providing a complementary interpretation of our shift-share variable:

$$Shiftshare_o = \underbrace{\sum_d \omega_{do0}}_{MigInc_{o0}} \times \underbrace{\frac{\sum_d \left(\omega_{do0} \tilde{\Delta} R_d \right)}{\sum_d \omega_{do0}}}_{Rshock_o} \tag{2.2}$$

$Shiftshare_o$ is the product of two terms. $MigInc_{o0}$ is pre-shock migrant income per capita in origin province o, across all migrant destinations. Provinces with higher $MigInc_{o0}$ have more migrant income per capita facing exchange rate risk (greater aggregate exposure to exchange rate shocks). $Rshock_o$ is the province-o weighted average exchange rate shock, where the weights are pre-shock shares of migrant income from each destination d. In Section 2.4 below, we emphasize that we derive causal identification solely from $Shiftshare_o$, not either of the component factors $MigInc_{o0}$ and $Rshock_o$ alone.

A key challenge is that the data needed to estimate exposure weights ω_{do0}, destination-d pre-shock migrant income per capita of province o, are not available in any Philippine Censuses or surveys. We estimate exposure weights ω_{do0} using two datasets from Philippine government agencies OWWA and POEA. The OWWA dataset contains the Philippine home address of individuals departing on overseas work contracts. The POEA dataset provides data on migrant income and occupation. Both the OWWA and POEA data include name, date of birth, destination, and gender. We match the two datasets to determine migrant origin province in the POEA database, and can then estimate ω_{do0}.[13]

[12]Borusyak et al. (2022a) call these terms "exposure shares", but we say "exposure weights" since they are not shares in our application. Because the sum of our ω_{do0} across destinations (within origins) is not one, we are in the "incomplete shares" case.

[13]We achieve a match rate of 95%. Further details of the matching process are in Appendix Section B.1.1.

Data for the exchange rate shock $\tilde{\Delta} R_d$ in $Shiftshare_o$ comes from Bloomberg LP. As we discuss in Subsection 2.4.2.1, our shift-share variable uses only the immediate, short-run change in exchange rates. We calculate the short-run exchange rate change, $\tilde{\Delta} R_d$, as the proportional change in the average exchange rate (foreign currency per PhP) from immediately before (mean from Jul 1996 - Jun 1997) to immediately after (mean from Sep 1997 - Oct 1998) the shock (e.g., a 10% appreciation of the foreign currency against the Philippine peso is 0.1).

2.3.2 Outcome Data

Provincial mean household income and expenditure per capita are available from the Family Income and Expenditure Survey (FIES), conducted every three years by the Philippine Statistics Authority (PSA). Each triennial FIES round samples roughly 40,000 households nationwide. We use up to twelve rounds of the FIES from 1985 to 2018 (inclusive), covering up to four pre-shock observations (prior to 1997), the "partially-treated" 1997 observation, and up to seven post-shock observations for each province.[14]

Key outcomes include migrant income, domestic income, and (their sum) global income per capita. We analyze these outcomes at the same triennial frequency as the FIES, the data source for domestic income. The POEA/OWWA contract data are available for fewer years, and also have missing data on migrant origin address in the early-to-mid 2000s (details in Appendix B.1), preventing us from calculating migrant income in 2000, 2003, and 2006. It is also not available after 2016. Analyses of migrant, domestic, and global income therefore involve fewer triennial periods: 1994, 1997, 2009, 2012, and 2015. Also in triennial periods, we examine secondary outcomes such as migrant contracts as share of province population (by occupation), and domestic income sub-components (wage, entrepreneurial, other). Income and expenditure outcomes are in 2010 real Philippine pesos (17.8 PhP/US$ PPP).

We also examine impacts on provincial educational attainment from six rounds of the Philippine Census of Population (1990, 1995, 2000, 2007, 2010, and 2015).

2.3.3 Import and Export Shift-Share Variables

Exchange rate shocks can potentially affect provincial outcomes through trade ties. An omitted variable concern arises if our migrant income shift-share measure is correlated with corresponding trade-based shift-share measures capturing potential impacts via imports and

[14] We exclude the partially-treated year 1997 from regression analyses, but include it in event-study analyses.

exports. We therefore construct import and export shift-share measure to assess the stability of our results to their inclusion.

The import and export shift-share variables are in the same spirit as our migrant income shift-share variable. The import and export shift-share variables exploit (pre-1997) variation in exchange rate shocks in import and export partners, in combination with province-level employment shares in import and export industries.

First, we compute the value of imports and exports between the Philippines and each partner country (destination) for each Standard International Trade Classification (SITC) good using COMTRADE data. We calculate these for the pre-Asian Financial Crisis period of 1990-1996. We aggregate the SITC goods to 36 ISIC industries to compute industry-level imports and exports between the Philippines and each partner country (destination).[15] Then, using the 1990 Population Census, we apportion the total industry-destination level import and export values to each province using the share of total Filipino workers in an industry that are in a given province. Summing up across industries yields province-destination level baseline import/export values. We divide this measure by province population to get a proxy for per capita import/export values between a given province and partner country.

Finally, we multiply this province-destination exposure measure by the destination exchange rate shocks, and sum over all destinations to get province-level shocks. Formally:

$$Shiftshare_o^m = \sum_d \underbrace{\frac{1}{Pop_o} \sum_j \frac{L_{jo}}{L_j} M_{jd}^m}_{\substack{\text{Per capita import/export} \\ \text{between } o \text{ and } d}} \tilde{\Delta} R_d \tag{2.3}$$

where $m \in \{import, export\}$ specifies the trade shock, o is province, d is destination (partner) country, and j is industry. M_{jd}^m is the total baseline value of industry j imports or exports between the Philippines and country d. L denotes the number of workers and Pop denotes population. $\tilde{\Delta} R_d$ is the exchange rate shock as before. This yields import and export shift-share variables that we include in the regression to gauge the robustness of the coefficient on the migrant income shift-share variable (in Panel D of all regression results to come).

2.4 Empirical Approach

We discuss the regression equation, causal identification, and temporal persistence of the shock measured by our shift-share variable.

[15] We match COMTRADE SITC data with ISIC revision 2 data using a crosswalk from World Integrated Trade Solution by the World Bank. Because the crosswalk is not complete, we manually match all remaining SITC products.

Table 2.2: Summary Statistics

	Mean	SD	10th P.	25th P.	Median	75th P.	90th P.	Obs.
Shock Variables								
Residualized $Shiftshare_o$	0.000	0.093	-0.105	-0.040	0.002	0.031	0.084	74
$MigInc_{o0}$	4.044	2.984	0.967	1.684	3.072	5.974	8.616	74
$Rshock_o$	0.415	0.040	0.371	0.389	0.412	0.436	0.454	74
Import Shock	10.673	9.180	2.766	4.148	7.661	12.864	24.678	74
Export Shock	10.432	10.057	2.786	4.626	6.801	12.739	21.966	74
Expenditure and Income								
Expenditure per Capita	29.074	10.525	18.220	22.041	26.939	33.557	42.329	887
Global Income per Capita	35.305	12.468	22.427	26.652	32.484	41.215	52.412	296
Domestic Income per Capita	30.699	10.618	20.007	23.453	28.570	35.151	44.949	296
Migrant Income per Capita	4.606	2.924	1.537	2.310	3.746	6.608	8.812	296
Education and Migration								
Share Primary School	0.789	0.114	0.638	0.719	0.799	0.880	0.927	444
Share Secondary School	0.486	0.146	0.291	0.374	0.490	0.580	0.689	444
Share College	0.133	0.046	0.082	0.098	0.126	0.158	0.191	444
Share College: Migrants	0.338	0.135	0.174	0.236	0.336	0.433	0.530	444
Migrant Share	0.013	0.009	0.003	0.006	0.011	0.018	0.025	444
Migrant Contracts (per 10,000 working age people)								
1st Quartile Education Occupations	94.191	71.725	22.301	44.736	82.824	120.183	178.979	296
2nd Quartile Education Occupations	8.694	6.616	1.730	3.760	6.886	12.455	16.924	296
3rd Quartile Education Occupations	24.690	19.297	5.942	12.679	19.967	34.584	47.180	296
4th Quartile Education Occupations	43.096	32.762	7.236	17.110	35.481	62.302	87.562	296
Baseline Province Controls								
Baseline Share Rural	0.643	0.193	0.337	0.564	0.696	0.761	0.819	74
Baseline Asset Index	-0.636	1.023	-1.576	-1.321	-0.966	-0.169	1.069	74
Baseline Total Income per Capita	29.914	10.333	20.504	23.191	27.803	32.582	46.112	74
Baseline Expenditure per Capita	24.368	7.891	16.416	19.454	22.683	26.817	35.265	74
Share of Workforce in Primary Sector	0.567	0.175	0.282	0.491	0.596	0.692	0.760	74
Share of Workforce in Industry	0.121	0.082	0.042	0.066	0.095	0.150	0.256	74
Share of Workforce in Service Sector	0.299	0.095	0.194	0.234	0.287	0.348	0.421	74
Share of Workforce in Financial Services	0.013	0.013	0.004	0.006	0.009	0.015	0.026	74
Baseline Destination Controls								
1995 GDP Per Capita	21.721	13.245	7.691	12.565	23.497	28.691	43.429	104
Average Contract Salary	329.291	258.947	108.387	108.387	166.838	669.068	708.831	104
Share of Contracts Professional	0.351	0.429	0.002	0.012	0.154	0.962	0.994	104
Share of Contracts Manufacturing	0.285	0.305	0.001	0.001	0.179	0.477	0.716	104
Share of all 1995 Contracts	0.126	0.098	0.011	0.024	0.108	0.192	0.299	104

Note: Unit of observation is 74 provinces (times periods as relevant) in all cases except bottom panel. For bottom panel, unit of observation is 104 migrant destination countries. Shock variables are constructed from POEA/OWWA dataset and other sources (see text). $MigInc_{o0}$ denotes pre-shock (1995) migrant income per capita. $Rshock_o$ denotes weighted-average exchange rate shock. Import and export shocks are as described in Section 2.3.3. Expenditure, total income, and domestic income data are from FIES. Migrant income is constructed from POEA/OWWA dataset and Philippine Census. Income and expenditure variables are in thousands of real 2010 Philippine pesos (17.8 PhP per PPP US$ in 2010). Periods for expenditure and total income are triennial, from 1985 to 2018 inclusive. (One observation, Rizal province in 1988, is missing due to loss of FIES data in a fire.) Periods for global, domestic, and migrant income data are 1994, 2009, 2012, and 2015. Shares of population by education level and share of population migrants are from Census (periods are 1990, 1995, 2000, 2007, 2010, 2015). Shares of population with primary, secondary, and college education are for those aged 20-64. "Share College: Migrants" is share of migrants reported in Census who have college or more education. Migrant contracts are from the POEA/OWWA dataset (periods are 1994, 2009, 2012, and 2015); working age defined as 20-64. Baseline province controls are from Census for share rural and asset index; and from FIES for total income and expenditure. Service sector excludes financial services (examined separately). Per capita GDP is from the World Development Indicators, in thousands of 1995 USD. Destination level contract controls are calculated from POEA/OWWA dataset.

2.4.1 Regression Equation

We estimate causal effects using the shift-share approach of Borusyak et al. (2022a). Our regression equation is:

$$
\begin{aligned}
y_{ot} = \alpha_o + \gamma_t + \beta_1(Shiftshare_o \times Post_t) \\
+ \boldsymbol{\lambda}'(MigInc_{o0} \times \mathbf{D}_t) + \boldsymbol{\phi}'(Rshock_o \times \mathbf{D}_t) + \boldsymbol{\delta}'(\mathbf{X}_{o0} \times Post_t) + \varepsilon_{ot}, \quad (2.4)
\end{aligned}
$$

y_{ot} is an outcome of interest for province o in period t. $Shiftshare_o$ is the shift-share variable, which is interacted with $Post_t$, an indicator for periods after 1997.[16] The coefficient β_1 is the coefficient of interest. Causal interpretation of β_1 exploits changes in migrant income per capita driven by the 1997 exchange rate shocks, as discussed in Subsection 2.4.2.1 below.

$MigInc_{o0}$ is pre-shock migrant income per capita in the province, and $Rshock_o$ is the province-o weighted-average exchange rate shock. Both these variables are interacted with a vector of period fixed effects $\mathbf{D}_t$.[17] Inclusion in the regression of $MigInc_{o0} \times \mathbf{D}_t$ and $Rshock_o \times \mathbf{D}_t$ accounts for changes from before to after the shock related to $MigInc_{o0}$ and $Rshock_o$. Identification of β_1 therefore derives solely from the interaction between $MigInc_{o0}$ and $Rshock_o$ embodied in $Shiftshare_o \times Post_t$.

$\mathbf{X}_{o0} \times Post_t$ is a vector of pre-shock destination characteristics and province-level characteristics interacted with the post-shock dummy. We discuss these further in Subsection 2.4.2.1. Province fixed effects α_o account for time-invariant differences across provinces. Period fixed effects γ_t account for common time effects. ε_{ot} is a mean-zero error term.

We do not impose the typical assumption of i.i.d. data. Our "shifts", the destination-d exchange rate shocks $\tilde{\Delta} R_d$, are common to provinces with similar exposure weights ω_{do0}. Borusyak et al. (2022a) and Adao et al. (2019b) demonstrate that conventional standard errors in shift-share designs are invalid due to likely correlation in residuals across observations with similar shock exposure. We report "exposure-robust" standard errors based on estimation of shock-level regressions following Borusyak et al. (2022a).

[16] While in many shift-share research designs the shift-share variable is used as an instrumental variable for a potentially-endogenous right-hand-side variable of interest, in our context we do not do so, and simply examine the "reduced form" impact of the shift-share variable. We take this approach due to likely violations of the IV exclusion restriction. Using $Shiftshare_o$ as an instrument for migrant income per capita, for example, would violate the IV exclusion restriction because the shock's effects operate not only via migrant income *per se*, but also via increased returns to migration. Perceived returns to education may then rise, driving education investments independently of effects due to migrant income shocks.

[17] Following Borusyak et al. (2022a), it is essential to interact the sum of exposure weights (which they call "sum of exposure shares") $MigInc_{o0}$ with period indicators in shift-share designs with incomplete shares and panel data. Time period fixed effects (the vector $\mathbf{D}_t$) alone will not isolate variation in the shock within periods. $MigInc_{o0} \times \mathbf{D}_t$ accounts for any time-period effects that vary according to $MigInc_{o0}$.

2.4.2 Causal Identification

We discuss assumptions required for causal identification, and empirical evidence supporting these assumptions.

2.4.2.1 Exogeneity of Exchange Rate Shocks

In the Borusyak et al. (2022a) shift-share approach, causal identification is based on exogeneity of the shifts (shocks), rather than on exogeneity of the shares. Our shifts are destination-d exchange rate shocks, $\tilde{\Delta} R_d$. The shares are province-o "exposure weights", ω_{do0}, for each destination.

Our identification assumption is therefore that the exchange rate shocks $\tilde{\Delta} R_d$ are as good as randomly assigned (conditional on destination-d-level controls). The exposure weights (shares) ω_{do0} can actually be endogenous.[18] An example of a failure of this assumption would be if a destination's exchange rate shock were correlated with the characteristics of Filipino migrant workers in the destination. For example, it would be a worry if baseline (pre-shock) migrant wages or education levels in a destination were associated with the destination's exchange rate shock.[19] Our estimate of β_1 in equation 2.4 could then be biased by any ongoing trends related to migrants' baseline characteristics.

Define the destination-d exchange rate shock immediately after the crisis as $\tilde{\Delta} R_d = \frac{R_{d,1998} - R_{d,1996}}{R_{d,1996}}$. $R_{d,1996}$ is the destination-d exchange rate (nominal Philippine pesos per destination-d currency unit) in the pre-period (twelve months leading up to June 1997), while $R_{d,1998}$ is the destination-d exchange rate in the immediate post-Crisis period (twelve months through October 1998). The exchange rate shock is thus a fractional change (e.g., a 10% appreciation is 0.1).

All components of the shift-share variable (equation 2.3) are from the pre-shock period, except for the post-shock exchange rate $R_{d,1998}$. Identification derives from the change in the destination-d exchange rate relative to its pre-shock level, $R_{d,1996}$.

It is plausible *a priori* that the exchange rate shocks are exogenous. The Asian Financial Crisis was unanticipated by global policy-makers and governments (Radelet and Sachs, 2000), so our estimates are unlikely to be clouded by anticipation of the shocks by households, firms, or officials in Philippine provinces (i.e., there are plausibly no effects of being treated in the future on outcomes in the pre-treatment period). While the real effects of the Crisis were short-lived (Park and Lee (2002) describes the "speedy V-shaped recovery"), the changes in

[18] In the Goldsmith-Pinkham et al. (2020) approach, the shares must be considered exogenous.

[19] Time trends in key outcomes such as migrant wages or employment may differ by baseline (pre-shock) values of the outcomes, for example if there are different growth rates across industries with different skill-intensities in production.

exchange rates were persistent.

Our shift-share variable exploits the fact that the Asian Financial Crisis was a surprise, using only the short-run (1997-1998) change in exchange rates immediately post-Crisis. We do not exploit further (post-1998) changes in exchange rates for identification. The short-run Crisis-induced exchange rate shocks are most plausibly exogenous. In the longer run, by contrast, the evolution of exchange rates may be endogenous to destination-country economic policies.

As it turns out, there is strong persistence of the short-run (1997-1998) exchange rate shocks over our entire two-decade study period. Destination-d 1997-1998 exchange rate shocks have strong predictive power for the long-run exchange rate up to 2018. We show this empirically in Subsection 2.4.4 below. By focusing on a shift-share variable defined with only the short-run 1997-1998 shocks, we estimate a reduced-form effect that includes any long-run exchange rate movements that are correlated with the short-run 1997-1998 exchange rate shocks, but that are not endogeneous to subsequent destination-level economic policies.

Since exogenous variation in this framework derives from the shifters (Borusyak et al., 2022a), we statistically show balance in these destination-specific exchange rate shocks. We run regressions at the level of all 104 migrant destinations. The dependent variable is the exchange rate shock, $\tilde{\Delta} R_d$, and the independent variables are pre-shock destination-d characteristics.[20]

The destination characteristics we examine are all pre-shock (1995). GDP per capita accounts for destination development status. Other independent variables are aspects of the destination's Philippine migrant flow. We account for the skill level of migrants going to particular destinations by, first, examining mean annual income per Philippine migrant in the destination. Second, we examine the share of Philippine migrants to the destination working in professional occupations (the highest-skilled occupation group), and separately the share of Philippine migrants to the destination working in manufacturing occupations (the intermediate-skilled group). We omit the lowest-skilled occupation group, services. In addition, we examine the share of all Philippine migrants going to the destination; this accounts for differences related to the aggregate size of the country as a migration destination. We also test the predictability of the exchange rate shocks with a sixth independent variable, the pre-shock (1994-1996) change in the exchange rate.[21] In a final regression we include all six independent variables.

Regression results in Appendix Table B.1 show no statistically significant relationships

[20]Following Borusyak et al. (2022a), observations in these regressions are weighted by the destination's average exposure weight ω_{do0} across provinces.

[21]Table 2.1 shows the change in the exchange rate in the pre-crisis period (1994-1996) alongside the change in the post-crisis period (1997-1998) for the top 20 destinations.

between pre-shock destination characteristics and the exchange rate shocks $\tilde{\Delta} R_d$. We reject joint significance of the right-hand-side variables in Column 7. These results provide support for the assumption that destination-d exchange rate shock can be considered as-good-as-randomly assigned.

While $\tilde{\Delta} R_d$ is balanced *vis-a-vis* these destination-level variables, inclusion of these controls can improve precision of estimates by absorbing residual variation. We therefore include these destination-level variables (interacted with the post-shock-period indicator) in the vector of controls $\mathbf{X}_{o0}$ in equation 2.4 (aggregated to the province level using exposure weights ω_{do0}, following Borusyak et al. (2022a)).

2.4.2.2 Exogeneity of Shift-Share Variable

Exogeneity of the exchange rate shocks should lead to exogeneity of our shift-share variable, $Shiftshare_o$. Concerns about causal identification arise if $Shiftshare_o$ is correlated with baseline (pre-shock) provincial characteristics (conditional on other right-hand-side variables in the regression). For example, provinces with lower baseline development status (income and expenditure per capita, rural share of population, etc.) could be on different time trends than other provinces.[22] If there are such differential time trends, and $Shiftshare_o$ is correlated with baseline (pre-shock) provincial development status, our estimate of β_1 in equation 2.4 would be biased. Thus it is important to control for potential differential time trends related to baseline development status of provinces.

As equation 2.2 shows, $Shiftshare_o$ can be written as the product of two terms. $MigInc_{o0}$ is migrant income per capita in province o in the pre-shock period. $Rshock_o$ is the province-o weighted average exchange rate shock. Table 2.2 shows $MigInc_{o0}$ has mean PhP 4,044 (standard deviation 2,984), while $Rshock_o$'s mean is 0.415 (standard deviation 0.040).

We take only $Shiftshare_o$ to be exogenous, not its component factors $MigInc_{o0}$ and $Rshock_o$. In regression equation 2.4, we achieve this by interacting $MigInc_{o0}$ and $Rshock_o$ with period fixed effects, which accounts for any changes over time that are correlated with these variables. Identification therefore comes only from $Shiftshare_o \times Post_t$.

The shift-share variable $Shiftshare_o$ is uncorrelated with pre-shock province characteristics, once $MigInc_{o0}$ and $Rshock_o$ are controlled for. This is apparent in Appendix Table B.2. There is no statistically significant relationship between $Shiftshare_o$ and pre-shock measures of provincial development. These results bolster confidence in the exogeneity of

[22] Initially-poorer provinces could be the beneficiaries of national government programs to improve education, promote small enterprises, improve infrastructure, etc., leading them to have more-positive time trends in development outcomes over our study period. The time trend could go in the opposite way, for example if agglomeration economies lead to higher growth rates in initially-richer provinces compared to initially-poorer ones.

$Shiftshare_o$ (after conditioning on $MigInc_{o0}$ and $Rshock_o$).

Because we only consider $Shiftshare_o$ exogenous when conditioning on $MigInc_{o0}$ and $Rshock_o$, we report in Table 2.2 the residualized $Shiftshare_o$ after partialling-out $MigInc_{o0}$ and $Rshock_o$. It has a mean of 0 and a standard deviation of 0.093. We will use this standard deviation of 0.093 in all discussions of magnitudes of effects below.

Figure 2.1 displays the spatial distribution of residualized $Shiftshare_o$ across provinces. The shock appears to be evenly distributed across the Philippines. All regions contain provinces with a range of shock values.

The pre-shock province-level characteristics examined in Appendix Table B.2 are also included in the control vector $\mathbf{X}_{o0}$ of regression equation 2.4. These controls capture changes over time that may be related to provincial pre-shock development. Inclusion of these controls can help improve precision by absorbing residual variation.

2.4.2.3 Falsification Tests

Following Borusyak et al. (2022a), we conduct a variety of falsification tests of the key assumption that the destination-d-level exchange rate shocks $\tilde{\Delta} R_d$ are as-good-as-random. Above, we showed that $\tilde{\Delta} R_d$ is uncorrelated with a variety of pre-shock destination characteristics (Section 2.4.2.1), and that the resulting shift-share variable $Shiftshare_o$ is conditionally uncorrelated with a set of pre-shock province characteristics (Section 2.4.2.2).

In addition, Borusyak et al. (2022a) also recommend conducting "pre-trend" analyses, testing whether changes in the outcome variable in the pre-shock period are correlated with the future value of shift-share variable. This is analogous to tests of parallel trends in difference-in-difference research designs. We present these in Section 2.5 (Appendix Table B.3) below. We find no evidence of that changes in any of our primary or secondary outcome variables in the pre-shock period are correlated with (future) $Shiftshare_o$. We also show event-study graphs of lead and lag coefficients of $Shiftshare_o$, building on regression equation 2.4 (Figure 2.3 and Appendix Figure B.8). These figures confirm the conclusion that pre-trends are uncorrelated with the future value of the shift-share variable.

2.4.3 Additional Threats to Identification

Impacts Thorough Trade. A key potential concern regarding whether the coefficient β_1 solely reflects changes in migrant income is that exchange rate shocks can also impact trade flows due to relative price changes. If migrant income shocks provinces face are correlated with such trade shocks, β_1 would be jointly capturing the impacts of trade shocks and migrant income shocks, complicating the interpretation of β_1. To provide direct evidence against this, we demonstrate the stability of our β_1 estimates to the inclusion of the import and export shift-share variables (discussed in Section 2.3.3) in the control vector $\mathbf{X}_{o0}$. Table B.4 demonstrates that this stability is plausible. The import and export shock variables are not correlated with the migrant income shock after controlling for $MigInc_{o0}$, $Rshock_o$, and the baseline control variables included in the main analyses (i.e. the variation that is relevant for our estimation).

We provide additional evidence in Section 2.5.3 that exports and FDI do not respond to the shocks of interest, and so do not appear to be mechanisms driving our findings.

Internal Migration. We also address the possibility of confounding changes in population composition. We examine the relationship between $Shiftshare_o$ and internal migration rates. Results are in Appendix Table B.5. We find no large or statistically significant impact on net internal migration. There is a small negative effect on outmigration, driven by young adults (aged 16-24), that cannot account for the impacts we document in our analyses. Changes in population composition due to internal migration appear to be a minor concern.

2.4.4 Persistence of Shock

We study the impact of changes in migrant income on long-run provincial outcomes, exploiting an exogenous shock measured by our shift-share variable. A key interpretive question is whether the shock is transitory or persistent.

We examine whether the shift-share variable's components – in equation 2.3, the exchange rate shocks $\tilde{\Delta} R_d$ (the "shifts") and the exposure weights ω_{do0} (the "shares") – show persistence over two decades post-1997. If both these components of the shift-share variable show persistence in the long run, the shock to migrant income would also be persistent.

We first examine persistence of the exchange rate shocks. Figure 2.2 shows nominal exchange rates (foreign currency units per PhP, normalized to 1 in 1996) for eight major Philippine migrant destinations. The year of the Asian Financial Crisis, 1997, is denoted by the vertical dashed line. The 1997 exchange rate shocks appear persistent, showing no apparent reversion to pre-shock levels.

Regression analyses confirm this conclusion. We run regressions at the level of 104 destinations, where the dependent variables are the change in the exchange rate from pre-Crisis to a certain post-Crisis year, and the right-hand side variable is the short-run (1997-1998) shock, $\tilde{\Delta} R_d$.[23] We present coefficient estimates on $\tilde{\Delta} R_d$ from seven different regressions, for different post-shock time periods, in Appendix Figure B.1a. Higher (more positive) coef-

[23] Observations are weighted by 1995 migrant income to that destination, following Borusyak et al. (2022a) for any destination-level regressions.

Figure 2.2: Exchange Rate Shocks Due to 1997 Asian Financial Crisis

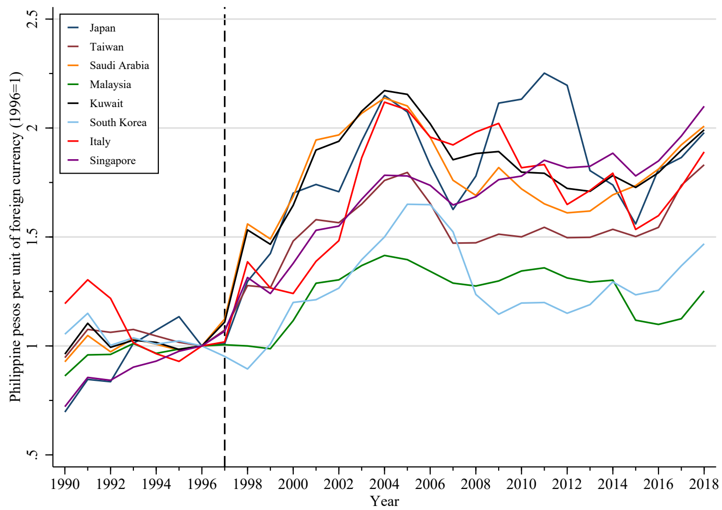

Notes: Data are from World Development Indicators. Annual average nominal exchange rates are in units of foreign currency per Philippine peso, normalized to 1 in 1996, for 8 large sources of international migrant income for Philippine provinces. Vertical dashed line indicates 1997 (year of the Asian Financial Crisis).

ficients indicate greater persistence, with a coefficient of 1 indicating complete persistence. Over nearly the entire study period, there is very strong persistence of the exchange rate shock. Point estimates are close to and statistically indistinguishable from 1 in nearly all post-shock periods. The only exceptions are 2009 and 2012, immediately following the 2007-2009 Great Recession, when the coefficients are closer to zero (very slightly negative in 2012), after which the coefficients rebound to levels near 1.[24]

Next, we analyze persistence of the the exposure weights ω_{dot}, migrant income per capita in destination-d/origin-o dyads. We create a dyad-level dataset with 7,696 observations (74 provinces times 104 destinations). For the post-shock periods for which we have migrant income data, we regress dyadic migrant income per capita in a post-shock year t (ω_{dot}) on dyadic migrant income per capita in 1995 (ω_{do0}), the pre-shock year in our shift-share variable. There is partial but substantial persistence over time in dyadic migrant income. Appendix Figure B.1b presents coefficients on ω_{do0} in the three regressions (for 2009, 2012, and 2015). The coefficients range in magnitude from 0.4 to 0.6. Each is statistically significantly different from zero (and from 1, indicating partial persistence).

In our theoretical framework, persistence in exposure weights ω_{dot} can stem from persistent dyad-specific migration costs, τ_{dot}, in equation B.2. While migrants adjust their post-1997

[24] As complementary support for the persistence of exchange rate movements, a Harris and Tzavalis (1999) test for a unit root in the 1990-2017 exchange rate panel data fails to reject the null of non-stationarity.

migration destinations in response to exchange rate changes, adjustment is only partial, due to networks facilitating migration (Munshi (2003b), Kleemans and Magruder (2019), Mahajan and Yang (2020b)), and (relatedly) information frictions in the international labor market (Shrestha and Yang (2019a), Shrestha (2020), Fernando and Singh (2021), Bazzi et al. (2021a)).

In sum, destination-level exchange rate shocks and dyadic migrant income per capita are highly persistent over two decades. The long-run impacts that we find result from an exogenous shock to migrant income (measured by the shift-share variable $Shiftshare_o$) that exhibits substantial persistence over time.

2.5 Empirical Results

We estimate impacts of the shift-share shock (β_1 in Equation 2.4) on a range of primary and secondary outcomes.

2.5.1 Domestic Income and Expenditure

We first examine impacts on key primary outcomes: province-level means of annual domestic income and expenditure per capita. We calculate these province-level outcomes from the FIES survey microdata.

"Domestic income" includes income from wages, entrepreneurial activity, and other sources, such as dividends, interest, and the imputed rental value of owned housing. We intend this outcome to capture household earnings in the *domestic* Philippine economy. This variable, therefore, does *not* include international migrant income (which in any case is not recorded in the survey), remittances, or other international income. (We calculate international migrant income using the migrant contract data and examine it in the next subsection.[25]) To avoid double-counting of earnings in the population, our measure of domestic income also excludes transfers from domestic sources and gifts from other households.

For expenditure per capita, we use the Philippine Statistical Authority's definition of "family expenditures": expenses or disbursements purely for personal consumption. This includes food, clothing, education, transport, communications, health, and utilities; con-

[25] By excluding international income sources from "domestic income", we are also excluding migrant remittances (which are not explicitly reported in the data; they are included in "overseas income"). There are concerns that migrant remittances are considerably under-reported in the FIES, because of the rise in electronic banking. Particularly since 2000, international migrants have been increasingly depositing their earnings directly into origin-household bank accounts. Comparison of remittance data from the World Bank, Philippine Central Bank, and the FIES suggests that households responding to the FIES may not consider funds deposited electronically into their bank accounts from overseas as remittances (Ducanes, 2010).

sumption from own production; and money payments made during the annual reference period for durable goods, furniture, and household repairs and maintenance.

The data are a panel of provinces observed every three years. There are four pre-shock observations (1985, 1988, 1991, and 1994) and seven post-shock observations (2000, 2003, 2006, 2009, 2012, 2015, and 2018) for each province. The 1997 observation is excluded because it is partially treated (the Asian Financial Crisis occurred in July 1997).

Results are in Table 2.3, columns 1-2. Each cell displays the coefficient β_1 on $Shiftshare_o \times Post_t$. We present estimates from regressions with different sets of pre-shock controls interacted with $Post_t$: destination controls only (Panel A), with additional province development status controls (Panel B), with additional province industrial structure controls (Panel C), and with additional import and export shift-share variables as controls (Panel D). All regression results tables will have this structure.

The shock has positive and statistically significant effects on both domestic income and expenditure per capita. Coefficient estimates in the domestic income regressions are stable across panels, and in Panel D the coefficient is statistically significantly different from zero at the 10% level. Coefficients in the expenditure regressions (column 2) are also stable across panels, and in Panel D the coefficient is statistically significantly different from zero at the 1% level.

The effects are large in magnitude. A one-standard-deviation shock (0.09) increases domestic income per capita by PhP1,349, and expenditure per capita by PhP1,224 (0.12 standard deviation in each case).

We also present event study diagrams illustrating dynamics of impacts, and testing for pre-trends. We estimate a modified Equation 2.4 in which we include the partially-treated year 1997 in the sample, and interact $Shiftshare_o$ with indicators for each time period. The 1994 interaction term is omitted as the reference point. We plot point estimates and 95% confidence intervals on $Shiftshare_o$ interacted with each period indicator. Results are presented in Figure 2.3a for expenditure and Figure 2.3b for domestic income. We do not observe differential positive pre-trends: for expenditure, pre-1997 coefficients are small and show no obvious trajectory. For domestic income, there is a slight negative trend from 1985-1991 and no trend in 1991-1994. There is also no large or statistically significant effect in 1997 for either outcome. For both outcomes, coefficients are positive and become larger over time after 1997. This increase in the magnitude of coefficients in the post-shock period is consistent with increases in domestic income per capita resulting from the gradual accumulation of human and physical capital over time.

We statistically confirm the absence of pre-trends with "placebo" regressions using the specification of equation 2.4, but for data in the pre-period (1985-1997 inclusive). We replace

Table 2.3: Effects of Migrant Income Shock on Global Income, Domestic Income, Migrant Income, and Expenditure per Capita

	Triennial: 1985 - 2018		1994, 2009, 2012, and 2015			
	(1) Domestic Income Per Capita	(2) Expenditure Per Capita	(3) Global Income Per Capita	(4) Domestic Income Per Capita	(5) Migrant Income Per Capita	(6) Expenditure Per Capita
Panel A. Destination controls only						
$Shiftshare_o \times$ Post	12.972	10.526	28.101	23.817	4.284	18.056
	(5.852)**	(4.045)***	(5.131)***	(4.122)***	(1.916)**	(4.626)***
Panel B. Additional province development status controls						
$Shiftshare_o \times$ Post	12.928	12.603	24.698	19.082	5.616	14.265
	(8.833)	(4.993)**	(7.926)***	(6.423)***	(2.453)**	(3.203)***
Panel C. Additional province industrial structure controls						
$Shiftshare_o \times$ Post	14.490	13.159	24.463	18.905	5.558	14.102
	(7.394)*	(4.726)***	(7.546)***	(5.982)***	(2.593)**	(3.473)***
Panel D. Additional import and export shift-share variables						
$Shiftshare_o \times$ Post	14.501	13.161	24.432	18.813	5.619	14.022
	(7.538)*	(4.909)***	(7.678)***	(6.656)***	(2.339)**	(3.729)***
Obs.	813	813	296	296	296	296
Dep. Var. Mean	29.885	28.975	35.305	30.699	4.606	30.181
Dep. Var. St. Dev.	10.908	10.505	12.468	10.618	2.924	10.623

Note: Unit of observation is the province-year. Domestic income and expenditure per capita are from Family Income and Expenditure Survey (FIES). Migrant income per capita is calculated from POEA/OWWA and Philippine Census data. Global income per capita is migrant income per capita plus domestic income per capita. Income and expenditure are in thousands of real 2010 Philippine pesos (17.8 PhP per PPP US$ in 2010). The year 1997 is dropped from the analysis as the exchange rate shock takes place in 1997. Outcome data are not available for one province (Rizal) in 1988 due to a fire that destroyed survey records. Destination pre-shock controls are (all for 1995): GDP per capita of the destination; mean annual income per Philippine migrant in the destination; share of Philippine migrants to the destination working in professional occupations (highest-skilled general occupational category); share of Philippine migrants to the destination working in manufacturing occupations (intermediate-skilled general occupational category; the lowest skilled general occupational category, services, is the omitted category); share of all Philippine migrants going to the destination. Destination controls are aggregated to the province level using Borusyak et al. (2022a) weights (province's pre-shock aggregate migrant income in the destination). Province development status pre-shock controls are as follows: share of households that are rural and household asset index (from 1990 Census); domestic income per capita and expenditure per capita (average across 1988/1991/1994 FIES). Province industrial structure pre-shock controls are as follows: share of workforce in primary sector, share of workforce in manufacturing, share of workforce in service sector, share of workforce in financial and business services (from 1990 Census). All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

Figure 2.3: Event Studies for Expenditure and Income per Capita

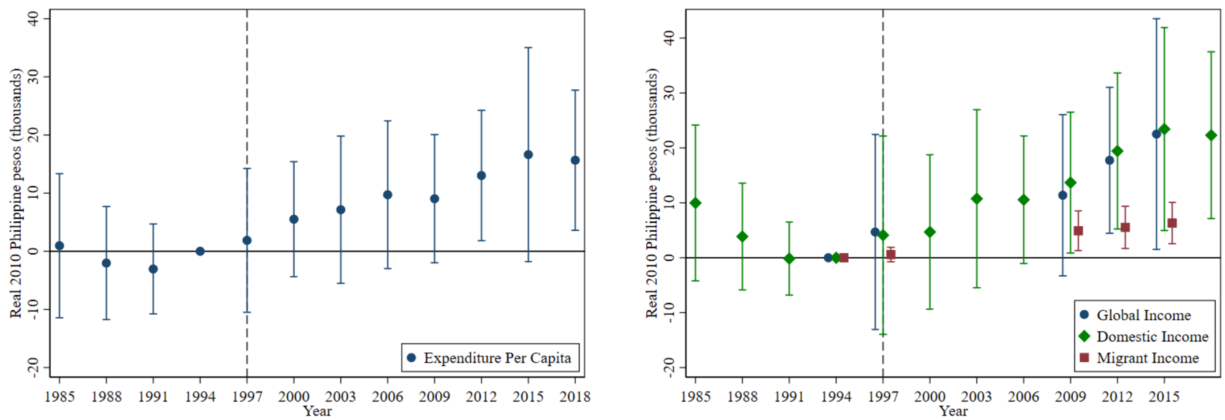

(a) Expenditure

(b) Global, Domestic, and Migrant Income

Note: Regressions modify Equation 2.4 to include interactions between $Shiftshare_o$ and indicator variables for each pre- and post-shock year. The 1994 interaction term is omitted as reference point. Specification corresponds to that of Table 2.3, Panel D (including province fixed effects, year fixed effects, and controls for differential trends with respect to pre-shock province characteristics, destination characteristics, and province import and export shift-share variables). Expenditure per capita includes food, education, durable goods, and housing, among other categories. Domestic income per capita includes earned income from wage and entrepreneurial activities, along with income from all other sources excluding transfers from abroad and domestic sources. Migrant income per capita is the sum of all income earned outside the Philippines by a province's migrants. Global income per capita is the sum of domestic and migrant income per capita. Outcomes are in real 2010 PhP (PhP17.8/US$ PPP). Observations are at the province-period level, and include each triennial period between 1985 and 2018 inclusive (when available); unlike in Table 2.3, we now include partially-treated year 1997 in the sample. 95% confidence intervals shown. Standard errors are clustered at the province level.

the indicator for the post-period, $Post_t$, with an indicator for a placebo post-period, 1994 and 1997. The years 1985, 1988, and 1991 are the placebo pre-period. Results are in the top panel of Appendix Table B.3, columns 1 and 2. The coefficients on $Shiftshare_o \times Post_t$ are small in magnitude and none are statistically significantly different from zero. These regressions confirm that there are no differential pre-trends.

2.5.2 Global, Domestic, and Migrant Income per Capita

We examine impacts on migrant income alongside impacts on domestic income. Migrant income is the sum of all income earned outside the Philippines by a province's international migrants. Domestic income is defined as in the above analysis: importantly, it excludes income from international sources. We also define "global income" as the sum of migrant income and domestic income.

Due to data constraints (see Section 2.3), we can only examine migrant and global income over five triennial periods: one pre-shock period (1994), one "partially-treated" period (1997), and three post-shock periods (2009, 2012, and 2015). In regression analyses, we exclude 1997, but include it in event-study analyses.

Regression results for global, domestic, and migrant income per capita are in columns 3-5 of Table 2.3. Within each Panel, the coefficient in column 3 is mechanically the sum of the corresponding coefficients in columns 4 and 5 (since global income is the sum of domestic and migrant income). The shock has positive and statistically significant effects on global, domestic, and migrant income per capita. Coefficient estimates are stable across regressions in Panels A , B, C, and D.

Impacts are large in magnitude. The coefficient estimate in column 3, Panel D indicates that each one-standard-deviation shock increases global income per capita by 2,272 pesos in 2009-2015 (0.18 standard deviation). Corresponding effect sizes for domestic income and migrant income per capita are 1,750 and 523 pesos, or 0.16 and 0.18 standard deviations respectively.

The coefficient estimate on migrant income (5.619) indicates that the initial shock to migrant income is substantially magnified over time: for each unit migrant income per capita shock (measured by our shift-share variable), migrant income per capita is over five times higher a decade later. We will turn shortly to the mechanisms behind this substantial magnification of the migrant income shock, examining the role of increases in migration rates, educational investments, and migrant skill levels.

To show the robustness of impacts on expenditure per capita, we also present regression estimates for this outcome in the restricted set of periods (1994, 2009, 2012, and 2015), in column 6. Point estimates and significance levels are very similar to the estimates of column 2 (which uses data from 1985-2018).

Figure 2.3b shows event study diagrams for migrant and global income per capita (along with domestic income results discussed above). There are no apparent pre-trends in the short 1994-1997 pre-shock period. The effects are positive in the 2009-2015 post-periods; point estimates are stable for migrant income, while global income point estimates are increasing. We also provide tests of the statistical significance of pre-trends in the bottom panel of Appendix Table B.3, columns 1 and 2. Pre-trend coefficients are small in magnitude and are not statistically significantly different from zero, confirming the absence of pre-trends.

2.5.3 Ruling Out Exports and FDI as Mechanisms

The stability of our estimates with the inclusion of import and export shift-share variables strongly suggests that the coefficient β_1 does not reflect the impacts of changes in trade flows due to the exchange rate shocks. Here we provide additional evidence pointing against exports driving the impacts we document. Further, we examine another potential mechanism, foreign direct investment (FDI), by testing whether aggregate FDI flows are affected by the same exchange rate shocks.

We first consider the value of manufactured exports per capita. We construct this outcome variable at the province-year level by aggregating firm survey microdata.[26] We estimate regression equation 2.4 where the dependent variable is in levels (PhP) and in inverse hyperbolic sine (IHS) transformation. We examine samples including all years (columns 1-2), as well as a restricted set of periods for "long run" results (1994-1996 vs. 2009-2015, columns 3-4). Results are in Appendix Table B.6. In no regression is there a large or statistically significant impact on manufactured exports.[27]

It is also of interest to examine agricultural exports, but no corresponding data exists for this outcome. We therefore examine agricultural income per capita, which should encompass any increase in agricultural exports. In Appendix Table B.7 we present regression estimates of equation 2.4 where the dependent variables are agricultural income per capita at the province-year level, in total as well as split into wage and non-wage (own production) income. We also show the impact on non-agricultural domestic income per capita for comparison. These outcomes come directly from the FIES data. The first four columns show results for the full set of triennial periods from 1985-2018, and the last four the periods for "long-run" results (1994 vs. 2009, 2012, and 2015).

The results in columns 1-3 and 5-7 reveal that there is no large or statistically significant impact on agricultural income (total, wage, and non-wage).[28] The impact of the shift-share shock on domestic income per capita appears to be entirely driven by the impact on non-agricultural income (columns 4 and 8). These results indicate that increases in agricultural export income (a subset of agricultural income) are unlikely to be driving the effects on domestic income.

Finally, we examine foreign direct investment (FDI) as a potential mechanism. Data on inward FDI from specific countries are not available at the province level, only at the national (Philippine) level (by year). We therefore run regressions analogous to Appendix

[26]These data are available in years 1994, 1996, 1998, 1999, 2006, 2009, 2010, 2012, 2013, 2014, and 2015. For further detail, see Appendix Section B.1.7.

[27]We note that even if we found impacts on manufactured exports, this would not necessarily mean our β_1 estimates are confounded by the impacts of exchange rate shocks on trade flows. An increase in domestic income due to migrant income shocks can, in principle, lead to increased exports. However, lack of an impact strongly suggests that shocks to exports are not a first order driver of our results.

[28]The standard deviation of the shift-share variable is is 0.093. The coefficients in both Appendix Tables B.6 and B.7 indicate that such a shock would have very small effects relative to the sample mean or standard deviation of either manufactured exports or agricultural income per capita.

Table B.1 (the tests for relationships between pre-shock overseas-destination characteristics and the exchange rate shocks), but this time in a panel context where the outcome variable is annual FDI flows to the Philippines from a particular country in a given year.[29]

The right-hand-side variable of interest is the exchange rate shock, $\tilde{\Delta} R_d$, interacted with a dummy for the post-shock period. The regression includes year and country fixed effects. We examine the full set of years (1996-2018, columns 1-2), the "long run" (comparing 1996 with 2009-2015, columns 3-4), as well as robustness to controls for overseas country characteristics (the same included in Table 2.3) in Panels A and B. Observations are weighted by the destination's average exposure weight ω_{do0} across provinces, following Borusyak et al. (2022a). This analysis tests whether the overseas-country-specific exchange rate shocks affect FDI flows to the Philippines *as a whole*. If no such relationship exists, it would be very unlikely that FDI flows to specific provinces are related to the shift-share shock. Results in Appendix Table B.8 indeed show no large or statistically significant relationship between FDI flows and the exchange rate shocks.[30]

Overall, these analyses provide no indication that exports or FDI are important mechanisms driving the causal effects emphasized in this paper.

2.5.4 Mechanisms

We now examine potential mechanisms through which these substantial increases in income take place. We examine educational investments, migrant skill levels and occupations, and domestic wage and entrepreneurial income.

2.5.4.1 Education

Relaxation of household liquidity constraints has been shown to lead to higher educational investments in the long run (Agte et al., 2022). Positive migrant income shocks could loosen such constraints on educational investments (Yang, 2008c; Gibson et al., 2011, 2014; Clemens and Tiongson, 2017; Theoharides, 2018b), and also change the expected return to education in the population at large.[31]

[29] These data are from the Philippine Statistics Authority. For further detail, see Appendix Section B.1.7.

[30] The standard deviation of the exchange rate shock, $\tilde{\Delta} R_d$, is 0.040. Appendix Table B.8's coefficients indicate that a shock of this magnitude would have very small effects relative to the mean or standard deviation of the outcome variable.

[31] Positive migrant income shocks could raise schooling investments overall if the return to education is perceived to rise (Batista et al., 2012; Docquier and Rapoport, 2012; Clemens and Tiongson, 2017; Shrestha, 2017; Theoharides, 2018b; Chand and Clemens, 2019; Khanna and Morales, 2023; Abarcar and Theoharides, 2022), but could reduce schooling investments if returns to education are seen to fall (McKenzie and Rapoport, 2011; de Brauw and Giles, 2017; Tang et al., 2022).

In Table 2.4 we present results from estimating regression equation 2.4 where the dependent variables are the share of the population having reached key threshold levels of education: primary (6 years of completed schooling), secondary (10 years), and college (14 years). Dependent variables are from the Philippine Census (pre-shock periods 1990 and 1995; post-shock periods 2000, 2007, 2010, and 2015). The positive shock to migrant income has positive and statistically significant effects on secondary and college (but not primary) completion rates.

Table 2.4: Effects of Migrant Income Shock on Education

	Share Completed:		
	(1) Primary School	(2) Secondary School	(3) College
Panel A. Destination controls only			
$Shiftshare_o$ × Post	-0.002	0.092	0.027
	(0.046)	(0.039)**	(0.030)
Panel B. Additional province development status controls			
$Shiftshare_o$ × Post	0.013	0.077	0.059
	(0.036)	(0.042)*	(0.028)**
Panel C. Additional province industrial structure controls			
$Shiftshare_o$ × Post	0.015	0.073	0.054
	(0.032)	(0.031)**	(0.019)***
Panel D. Additional import and export shift-share variables			
$Shiftshare_o$ × Post	0.014	0.073	0.054
	(0.042)	(0.022)***	(0.018)***
Obs.	444	444	444
Dep. Var. Mean	0.789	0.486	0.133
Dep. Var. St. Dev.	0.114	0.146	0.046

Note: Unit of observation is the province-year. Analysis uses Census data; periods are 1990, 1995, 2000, 2007, 2010, and 2015. Dependent variables are share of population (aged 20-64) who have completed primary, secondary (high school), and college education. Primary school, secondary school, and college completion is defined as having completed at least 6, 10, and 14 years of schooling respectively. For list of destination and provincial controls, see Table 2.3. All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

Coefficient estimates in columns 2 and 3 indicate that a one-standard-deviation migrant income shock causes 0.68 percentage points higher secondary completion, and 0.51 percentage points higher college completion. Point estimates in those regressions are relatively stable across sets of controls and statistically significantly different from zero at the 1% level in

Panel D.[32]

These educational responses to the shock are plausible in magnitude. We gauge magnitude plausibility by examining the extent to which the increases in education we document are associated with increases in household income, since loosened financing constraints are likely a key reason behind the increase in education. Our regression results, comparing Panel D of Table 2.3 (col 3) with Table 2.4 (column 3) indicate that about 4,524 pesos higher global income is associated with 0.01 higher college completion.[33]

How does this relationship between increased income and increased education compare to relationships seen in cross-sectional data in the pre-period? The cross-sectional relationship between global income and share skilled in the population in the pre-period (1994 for income and 1995 for education) indicates that each 0.01 higher college completion is associated with about 3,500 pesos more in provincial global income per capita. While this is not a causal effect, it is a reasonable point of comparison. The education response we estimate is slightly smaller: 4,524 PhP is "needed" to generate the same increase in college completion.

2.5.4.2 Migrant Skills and Occupations

The increase in education in the population may also raise migrant workers' skill levels. We first examine whether the shocks to migrant income have a causal impact on the share of migrants who are skilled, defined as having at least college (14 years) education. This outcome is available for international migrants in the Philippine Census. Periods included in the regression are the Census years 1990, 1995, 2000, 2007, 2010, and 2015.

In column 1 of Table 2.5, we report results from estimating equation 2.4 where the dependent variable is the share of international migrants who are skilled. There is a substantial positive effect that is stable across panels with different sets of controls. The coefficient in Panel D is statistically significantly different from zero at the 1% level. A one-standard-deviation higher shock leads to 1.8 percentage points higher share of migrants who are skilled (0.14 standard deviations).[34]

Is this increase in migrant educational levels associated with working in higher-skilled jobs? We examine impacts on the propensity to enter skilled international migrant work. These analyses require the migrant contract data, so the periods included in the regression

[32]Falsification tests in Appendix Table B.3 (middle panel, columns 1-3) and event-study graphs of lead and lag coefficients of $Shiftshare_o$ in Appendix Figure B.8, subfigure (b), confirm the absence of pre-trends for these education outcomes.

[33]Note of course that the increase in education investments due to the shock could also be driven in part by perceived changes in the return to education, not only by loosened financing constraints.

[34]For this outcome, there is no evidence of pre-trends in Appendix Table B.3 (middle panel, column 4) or in Appendix Figure B.8, subfigure (c).

Table 2.5: Effects of Migrant Income Shock on Contract Types and Migrant Skill

	Census	Contracts per 10,000 Working Age People			
	(1) Share Skilled Migrants	(2) 1st Qtile Education	(3) 2nd Qtile Education	(4) 3rd Qtile Education	(5) 4th Qtile Education
Panel A. Destination controls only					
$Shiftshare_o \times$ Post	0.165	16.546	3.249	64.683	57.422
	(0.052)***	(65.065)	(7.237)	(26.474)**	(15.382)***
Panel B. Additional province development status controls					
$Shiftshare_o \times$ Post	0.210	5.968	-0.509	55.807	28.393
	(0.061)***	(72.066)	(8.414)	(26.280)**	(17.146)*
Panel C. Additional province industrial structure controls					
$Shiftshare_o \times$ Post	0.196	1.044	-1.567	46.026	19.841
	(0.059)***	(73.254)	(8.873)	(21.630)**	(18.730)
Panel D. Additional import and export shift-share variables					
$Shiftshare_o \times$ Post	0.196	2.366	-1.591	46.569	20.207
	(0.059)***	(66.661)	(8.054)	(17.097)***	(19.542)
Obs.	444	296	296	296	296
Dep. Var. Mean	0.338	94.191	8.694	24.690	43.096
Dep. Var. St. Dev.	0.135	71.725	6.616	19.297	32.762

Note: Unit of observation is the province-year. Share of migrant workers who are skilled is from the Census (periods are 1990, 1995, 2000, 2007, 2010, and 2015). Skilled is defined as completing 14 years of education, which corresponds to finishing a college degree. Migrant contract variables are calculated from POEA/OWWA data (periods are 1994, 2009, 2012, and 2015). Outcome variables in columns 2-5 are migrant contracts (per 10,000 working age population) in occupations in the 1st (lowest) through 4th (highest) quartiles of migrant years of education. For list of destination and provincial controls, see Table 2.3. All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

are 1994, 2009, 2012, and 2015 (as in Table 2.3, columns 3-6). The dependent variable is migrant contracts per 10,000 working age (age 20-64) population.

We estimate equation 2.4 for migrant contracts in four quartiles of occupations, ordered from lowest (1st quartile) to highest (4th quartile) education levels.[35] Results are in columns 2-5 of Table 2.5. There are positive effects on new international migration in the two highest-education quartiles of occupations, but not for the bottom two quartiles. The coefficient is largest and statistically significant in Panel D for the 3rd (second to highest) quartile, while that on the 4th (top) quartile is also positive but not significantly different from zero.[36]

In sum, migrant income shocks increase the share of migrant workers who are skilled (have college or more education), as well as migrant flows in higher-education occupations. These effects are likely to be mechanisms leading to the substantial gains in income over the long run.

2.5.4.3 Entrepreneurial, Wage, and Other Domestic Income Sources

We now examine impacts on sub-types of domestic income. Table 2.6 presents regression results from estimating Equation 2.4 where dependent variables are domestic wage income, entrepreneurial and rental income, and other income per capita. Wage income is compensation (cash or in-kind) from regular or seasonal work. Entrepreneurial and rental income is from any entrepreneurial activity (such as poultry/livestock raising, wholesale/retail, transportation services, and rental of land/property). Other income includes pensions, interest, dividends, and other sources.

The shock led to increases in both wage income as well as entrepreneurial and rental income. Coefficient estimates for both these outcomes are robust to the set of controls. They are statistically significantly different from zero at conventional levels in Panel D, and similar to one another in magnitude. By contrast, there is no robust evidence that "other" income is a major part of the increase in domestic income. The positive impact

[35] The 4th (top) quartile (mean 14.4 years of schooling) includes engineers, medical professionals, and teachers. The 3rd quartile (mean 12.9 years of schooling) includes caregivers, restaurant workers, and performing artists. The 2nd quartile (mean 12.7 years of schooling) includes laborers and production workers. The 1st (bottom) quartile (mean 12.3 years of schooling) includes household workers (maids) and construction workers. Years of education data refer to 1992-1996 (pre-shock) contracts. The contract data do not include migrant worker education, so we calculate mean years of education in 80 detailed migrant occupations in the 1992-2003 Survey of Overseas Filipinos (SOF). We then assign the mean years of education for the occupation from the SOF to each migrant working in the occupation in the contract data. Then, we calculate mean migrant education within quartiles of the contract data. Quartiles are somewhat uneven in size due to lumpiness in the distribution of contracts across occupations.

[36] For these outcomes, we examine pre-trends in Appendix Table B.3 (bottom panel, columns 3-6) and in Appendix Figure B.8, subfigure (d). None of the coefficients in the pre-trend regressions are statistically significantly different from zero. The coefficient for the 1st (lowest-education) quartile is large in magnitude, suggestive of a differential positive pre-trend for that outcome (but note we report no effect on that outcome in Table 2.5). For the more-skilled (3rd and 4th) quartiles, the coefficients in the pre-trend tests are not negligible, amounting to about two-thirds the magnitude of the corresponding coefficients in Table 2.5. Overall, we view these tests as providing modest (but not overwhelming) support for the absence of pre-trends for these outcomes representing migration in high-skilled occupations.

Table 2.6: Effects of Migrant Income Shock on Components of Domestic Income

	Domestic Income Components:		
	(1) Wage Income	(2) Entrepreneurial and Rental Income	(3) Other Income
Panel A. Destination controls only			
$Shiftshare_o \times$ Post	10.022	9.741	4.054
	(3.081)***	(1.295)***	(2.122)*
Panel B. Additional province development status controls			
$Shiftshare_o \times$ Post	9.853	8.289	0.940
	(4.507)**	(1.991)***	(1.954)
Panel C. Additional province industrial structure controls			
$Shiftshare_o \times$ Post	9.733	7.881	1.291
	(3.690)***	(1.487)***	(2.160)
Panel D. Additional import and export shift-share variables			
$Shiftshare_o \times$ Post	9.691	7.866	1.256
	(3.730)***	(1.702)***	(2.481)
Obs.	296	296	296
Dep. Var. Mean	15.110	10.155	5.434
Dep. Var. St. Dev.	7.779	3.311	2.414

Note: Unit of observation is the province-year. Data from the Family Income and Expenditure Survey (FIES); periods are 1994, 2009, 2012, and 2015. For list of destination and provincial controls, see Table 2.3. All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

on wage income and on entrepreneurial and rental income are likely to reflect higher levels of education in the population, as well as increased capital investment in enterprises (both within and outside the household). We explore this further in Section 2.6 below.

2.6 Model-Based Quantification and Discussion of Magnitudes

We now provide further insight into mechanisms and magnitudes of the results thus far. First, we outline a theoretical framework to shed additional light on the long-run effects on global income and its components, migrant and domestic income. We take a simple model-based approach to quantifying the contribution of educational investments to the long-run income gains. The theoretical framework derives changes in skill shares, migration flows, migrant income, and domestic income as a function of the shift-share variable. In addition, the model allows us to shed light on whether the magnitude of the effect on migrant income per capita in the long run is explicable. We summarize this model-based quantification here. Appendix Section B.2 contains the full details, derivations and calculations underlying the model. It also presents validation tests that show our simple and tractable framework does a good job of predicting changes in migration rates and various sources of income.

Figure 2.4: Stylized Overview of Possible Channels

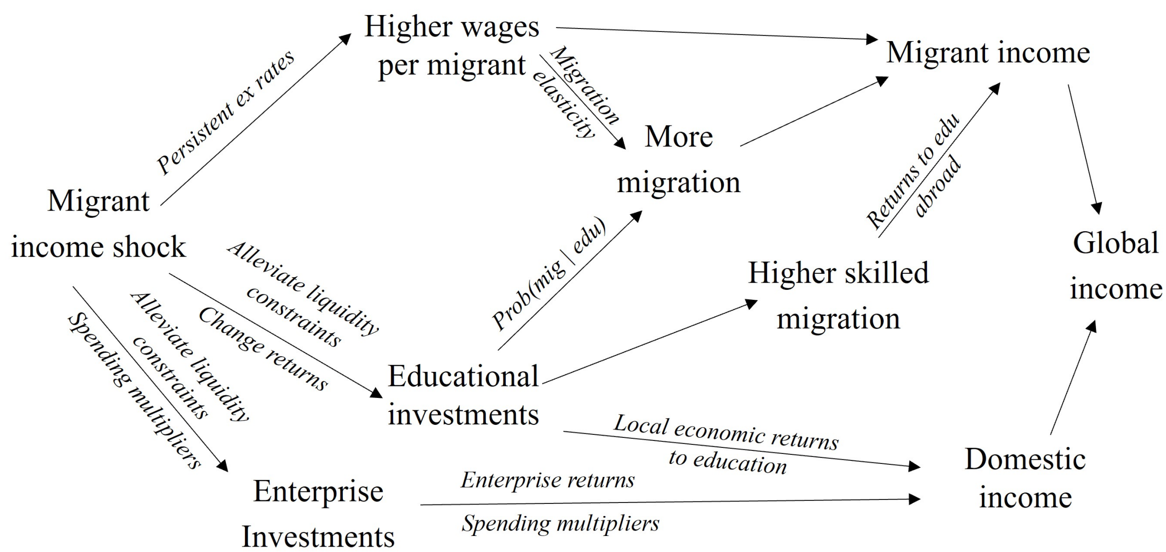

Note: Overview of modelled channels via which the migrant income shock affects global income. Details in Appendix B.2.

In Figure 2.4 we present a stylized diagram to describe the various channels in the model through which the migrant income shock may affect global income. The persistent migrant income shock drives higher wages per migrant; which in turn may lead to more migration and migrant income. The initial shock may also be invested in education, which may lead to more migration (as the skilled are more likely to migrate) in better-paying skilled jobs, again raising migrant income. The investments in education also drive increases in domestic earnings back home. If this overall high persistent migrant income is invested in domestic enterprises or drives local consumption spending demand, it may also raise domestic earnings. We provide full details of the model in Appendix Section B.2.

2.6.1 Contribution of the Education Channel

The long-run impact of the migrant income shock may be partly due to increased educational investments. First, skilled workers earn more. Furthermore, better-educated individuals have higher migration rates, and better-educated migrants work in higher-skilled jobs overseas. We quantify the contribution of educational investments in the long-run changes in both migrant and domestic income.

The college completion regression in Table 2.4 provides the estimate of the educational investment response to the shock. To estimate the contribution of educational investments to the income gains, we first multiply each province's specific value of the shift-share variable by the regression coefficient (0.054) in Panel D, column 3 of Table 2.4 to estimate the change in the province's population share skilled. Then we estimate how migration (to different destinations, as well as remaining at origin) would change in response to the change in the population skill composition, presuming the same dyadic migration probabilities by skill (the probability someone with skill s migrates from origin o to destination d) from the pre-shock period (1995). That is, to estimate the changes in migration flows to the various destinations, we first take the difference between skill groups in the baseline proclivity to migrate to various destinations, and multiply this difference by the change in the share skilled.

Then, we calculate how both migrant and domestic income would change in response to such migration changes, presuming the same dyadic skill premium (difference in skilled vs. unskilled income, in origin-destination dyads) from the pre-shock period. That is, we take the baseline skill premia, both for domestic and for migrant income, and multiply it by the change in share skilled to predict the education-driven change in incomes.

This calculation provides us with estimates of the change in migrant and domestic income per capita due to the education channel. We estimate that the education channel explains 24.1% of the increase in migrant income, and 22.9% of the increase in domestic income. Global income is the sum of migrant and domestic income; the implied share of global income explained by increased education is 23.2%. In sum, the increases in education induced by the exogenous increase in migrant income account for roughly one-fourth of long-run income gains.

2.6.2 Explaining Impact on Migrant Income

We also use the model to explain the large increase in migrant income, relative to the initial migrant income shock measured by the shift-share variable (the coefficient estimate of 5.619 in Table 2.3's migrant income regression). As discussed above, 24.1% of the increase in migrant income is explained by increased educational attainment. We seek to explain the remaining three-fourths of the migrant income increase. Additional mechanisms leading to further migrant income gains include the exchange rate shocks themselves, as well as changes in migration flows across destinations.

We first estimate changes in migration flows. Destination exchange rate shocks could change migration decisions, contributing to the eventual changes in long-run migrant income. In our gravity equation, the Fréchet parameter θ is the elasticity of migrant flows (from origin-o to destination-d) with respect to destination wages. This determines subsequent location choices and migrant income. Higher θ means that migration flows, and thereby migrant income, respond more to exchange rate shocks. We use the exchange rate shocks to estimate θ in Appendix B.2.4 using a Poisson pseudo-maximum likelihood (PPML) estimator (as many origin-destination dyads have zero flows). This yields an estimate of 3.42, which we use along with the actual exchange rate shocks to predict changes in migration in origin-destination dyads.[37]

We then calculate the change in total migrant income resulting from all dyadic (origin-destination) changes in migration flows, by skill, along with changes in destination exchange rates. We presume that skill-specific migrant wages (in destination currency) in each destination are fixed at pre-shock levels, so that changes in migrant income are driven only by exchange rate shocks and changes in migration flows. We estimate that these factors explain an additional 74.7% of the change in migrant income. This is on top of the 24.1% of the increase in migrant income attributed to education investments. The modeled components therefore explain almost all (98.8%) of the increase in migrant income.

In sum, the model accounts for the entire magnitude of the effect on migrant income. The five-fold magnification of the initial migrant income shock is fully explained by the combination of increased education, persistent exchange rate shocks, and changes in migration across destinations.

2.6.3 Explaining Impact on Domestic Income

We investigate assumptions needed to explain the magnitude of the impact on domestic income per capita. The coefficient on the shift-share variable in the domestic income per capita regression of Table 2.3, Panel D, column 4 indicates that a PhP 1 migrant income shock leads to a PhP 18.81 increase in long-run domestic income. 22.9% of this increase is attributable to the increases in education investments (see Subsection 2.6.1). This leaves

[37] We account for "indirect resorting": potential migrants simultaneously consider the full set of exchange rate changes in migration decisions, rather than simply choosing between migrating to specific destination-d or remaining at origin. For example, if Japan's exchange rate appreciates, while Malaysia's depreciates, migration to Malaysia will fall, but some individuals deterred from Malaysian migration will migrate to Japan instead of not migrating.

PhP 14.5 to be explained. We consider two mechanisms that could explain this remainder: a demand multiplier, and investments in domestic enterprises.

Recent studies have estimated large demand multipliers in low-income contexts. Egger et al. (2022) estimate a multiplier of 2.5 in response to cash transfers in Kenya. The multiplier due to a credit supply shock in India is 2.9 (Breza and Kinnan, 2021). We consider how much of our effect on domestic income could be explained by such multipliers. In our context, multipliers operate on the portion of migrant income sent back to origin provinces. The coefficient estimate in the migrant income regression of Table 2.3, Panel D indicates that the multiplier would operate on the portion of the 5.619 increase in migrant income per capita that is sent back to origin provinces. Assuming 70% of the migrant income returns to the local economy, that coefficient and a multiplier of 2.9 implies an increase in domestic income per capita of 11.28 PhP (5.558 x 0.7 x 2.9). A simple demand multiplier thus explains 77.8% of the remaining 14.5 PhP.

We now consider an additional contributor to the increase in domestic income: migrant income could alleviate constraints on capital investments. The migrant income shock was not a one-time windfall, but was sustained and grew over time, and so likely led to a sustained increase in capital accumulation. It is widely recognized that household enterprises and firms face binding constraints on capital investment (Karlan and Morduch, 2010), and that when such constraints are loosened, firms have high rates of return on investment. For example, de Mel et al. (2008) estimate a rate of return to Sri Lankan microenterprise investments from randomly-assigned capital investments of 5% per month (80% per year).[38] Such returns likely explain part of the increases in wage and entrepreneurial incomes we document in Table 2.6.

We examine whether our domestic income results can be generated in a stylized framework in which a portion of the exogenous increase in migrant income is devoted to capital accumulation in productive enterprises, and in which a demand multiplier also operates. We summarize the framework here; details are in Appendix Section B.2.7.1.

We trace the dynamics of domestic income per capita following the initial shift-share shock. Shock-induced migrant income per capita grows over time, reaching the amounts reflected in the event-study coefficients for migrant income per capita in Figure 2.3. In each post-shock year, a portion of shock-induced higher migrant income returns to origin provinces. Migrant income returned to origin economies generates an aggregate demand multiplier. In every period, households save a portion of shock-induced higher incomes,

[38]Similarly high returns are found by Banerjee and Duflo (2014), Hussam et al. (2022), and Cai and Szeidl (2022). In the Philippines, Edmonds and Theoharides (2020) find a rate of return of 27%, 18 months after a productive asset transfer (although Karlan and Zinman (2018) find limited savings constraints in the Philippines).

investing them in enterprises and firms.[39] We assume relatively high initial rates of return on investment (but not as high as the findings of de Mel et al. (2008)), which decline over time as the initial low-hanging investment fruits are exhausted. Higher incomes induced by these capital investments also generate a multiplier.

In Appendix Figure B.7a, we display the shock-induced domestic income of the model between 1998 and 2015, for three values of the share of migrant income spent at origin, α. With α=0.7, a PhP 1 initial migrant income shock becomes PhP 16.7 of domestic income by the year 2015. In Appendix Figure B.7b, we set α=0.7, and vary the initial rate of return on investment and trace the shock-induced domestic income in 2015. Our estimates range from 13.4 for a rate of return of 0.05, to 20.5 when the rate of return starts at 0.8 (the estimate of de Mel et al. (2008)).

We view this calculation primarily as a sanity check, demonstrating that a set of reasonable assumptions can generate the observed long-run impact on domestic income per capita. The framework does not incorporate all possible channels through which the effect on domestic income may arise. Importantly, we do not model potential escapes from poverty traps, such as those due to investment indivisibilities (Ghatak, 2015; Balboni et al., 2021; Kaboski et al., 2022). Considering escapes from poverty traps would make it even easier to explain the magnitude of the long-run effect on domestic income.

2.7 Conclusion

We study the long-run consequences of persistent increases in international migrant income for migrant-origin regions. We find that the vast majority of income gains are from *domestic* (origin-area) sources; gains in international migrant income, while also substantial, account for only a minority of gains. In addition, model-based estimates suggest that about one-fourth of the income gains (both domestic and international) are due to increased educational investments.

Our findings suggest that migration policy should be an important part of the development policy toolkit. Our results shed light on the impacts of policies – in both origin and destination countries – that affect current international migrant income as well as opportunities to earn such income in the future. Origin-country policies include efforts to facilitate international labor migration, as well as regulation to reduce market power of international labor market intermediaries (ensuring migrants retain more of their income gains). They might also include origin-country educational policies that raise population skill levels and

[39] We set the savings rate to 0.35, which implies a Keynesian multiplier of 2.86 (comparable to the 2.9 estimate in Breza and Kinnan (2021)).

make citizens more competitive for international jobs. Destination country policies include increases in legal immigration opportunities, enforcement against undocumented immigrants, and labor market policies that affect immigrants' ability to work legally. Our findings also have relevance for exchange rate policy in developing countries, highlighting that migrant-origin-currency devaluations can have positive long-run effects by raising current migrant income and returns to migration in migrant-sending areas.

There are also implications for how we think about overseas development assistance (foreign aid). We find that improvements in migrant income have substantial positive impacts on development of the *domestic* economy of migrant *origin* areas. Development agencies could consider supplementing traditional foreign aid with programs that facilitate international labor migration (Clemens, 2010; Clemens and Pritchett, 2013; World Bank, 2018a; Nunn, 2019).

CHAPTER 3

On the Allocation and Impacts of Managerial Training

with Achyuta Adhvaryu and Anant Nyshadham

3.0 Abstract

We study the allocation and productivity consequences of managerial training via a randomized controlled trial among production line supervisors in a large ready-made garment firm. We designed a program using practices identified as productive in Adhvaryu et al. (2022d), and asked middle managers – who are directly above production line supervisors in the hierarchy – to recommend which of the supervisors they manage should be prioritized for training. We then randomized access to the program within these recommendation rankings. Productivity on lines managed by treated supervisors increased by 6-7% relative to control, but these gains exhibit substantial heterogeneity across middle manager ranking categories. Highly recommended supervisors experienced no productivity gains; the average treatment effect of training is driven entirely by low-recommendation supervisors. This was not due to a lack of information about baseline skills or about who would gain the most, nor to discrimination or favoritism along observable dimensions. Instead, consistent with the fact that supervisor turnover has large personal costs for middle managers in terms of labor substitution and onboarding, middle managers prioritized the retention impacts of training. Treated supervisors were 14% less likely to quit than controls over the study period, and this gain was driven by highly recommended supervisors. Heterogeneous returns and the unproductive allocation of costly training can thus help explain underinvestment in attenuating persistent within-firm gaps in managerial quality.

3.1 Introduction

Variation in the quality of managers contributes to the vast dispersion in productivity across countries and firms (Bloom and Van Reenen, 2007, 2011; Bloom et al., 2016; Bandiera et al., 2020), and even across teams and workers within firms (Frederiksen et al., 2020; Bertrand and Schoar, 2003; Lazear et al., 2015). Many skills and practices have been identified as key to the productive value of managerial quality – for example, monitoring and dynamic task allocation (Adhvaryu et al., 2022a,d), cultivating good relationships with team members (Hoffman and Tadelis, 2021; Alan et al., 2023), and motivating and eliciting performance from workers (Frederiksen et al., 2020). Given their importance, why, then, do managerial skill gaps persist? In particular, what prevents training – which is perhaps the most common lever used by organizations of all kinds to upgrade managerial practices – from successfully closing these gaps?

We contend that the persistence of managerial skill gaps is due in part to the potential for misallocation of – and consequently low perceived returns to – training within the firm. Training is an endogenously allocated investment within organizations. *Who* allocates training – and how closely the objectives of these agents are aligned with the firm's – is thus of paramount importance. Even if the true average treatment effect of managerial training on productivity is large for the population of managers in a firm, realized impacts could be small if returns are heterogeneous enough and if training tends to be allocated to the "wrong" managers. This concern is perhaps most salient when firms pilot costly investments on a subset of workers or teams to inform full-scale adoption decisions. In these instances, middle managers, who work most closely with eligible candidates, often play a critical role in nominating pilot participants.[1]

This practice harkens to the literature on the trade-offs of decentralization of decision-making within organizations. Having multiple layers of management can be valuable (Caliendo and Rossi-Hansberg, 2012; Caliendo et al., 2015, 2020), particularly if it is possible to delegate some responsibilities and decisions to lower levels of the hierarchy (Bloom and Van Reenen, 2011; Bloom et al., 2014; Aghion et al., 2021). The argument is that middle managers may have private information or specialized understanding that makes them better equipped to make optimal resource allocation decisions. On the other hand, the classic trade-off is that this decentralization creates a principal-agent structure in which the middle manager will act according to her own incentives, which may not perfectly align with those of the organization, and that limited information at the top may make achieving first-best

[1] This was indeed the modal approach employed by our partner firm, and the primary motivation for our study design.

decision-making impossible (Atkin et al., 2017; Acemoglu et al., 2007; Aghion et al., 2014; Haegele, 2022b).

In this study, we ask: what are the impacts of training low-level managers in critical soft skills? And would middle managers allocate training to maximize productivity returns as the firm would want, or is their decision-making driven by other considerations? We seek to answer these questions via a randomized controlled trial in which we designed and implemented a management training program for production line supervisors in Indian garment factories. The program curriculum, which was designed by the authors along with the firm partner, emphasizes particular soft skills – such as communication, planning and organization, problem-solving, and motivation of workers – that were found to be strongly related to productivity in a nearly identical setting by Adhvaryu et al. (2022d). To study questions regarding allocation, we asked middle managers to rank the supervisors in their charge on who should be prioritized for training. We then randomized access to the training within these rankings, allowing us to recover average treatment effects as well as to study whether middle managers did indeed recommend the supervisors who ultimately gained the most from training.

We find that line supervisors gained substantial knowledge from the training, with test scores of treatment supervisors increasing by 40 to 100% as compared to control supervisors who exhibited no significant gains as expected. Productivity of teams managed by trained supervisors increased substantially and persistently on average compared to controls – by 7.3 percent during training and 5.8 percent over the six months after training completion. However, these average gains belie substantial heterogeneity across middle manager ranking categories. Line supervisors recommended highly by middle managers to receive the training actually experienced no significant productivity gains; the average treatment effect is driven entirely by low-recommendation supervisors.

This striking pattern of results is not due to middle managers lacking information about their supervisors' baseline skills, or about who might benefit most from training. It is also not due to discrimination or favoritism on observable dimensions. We learned from conversations with middle managers as well as senior management that while middle managers are indeed incentivized to achieve high productivity (through performance-based rewards), they also face large personal effort costs from supervisor turnover. In particular, middle managers are charged with the training and onboarding of new supervisors and performing supervisory duties in the interim, all of which involves substantial effort in addition to their other day-to-day responsibilities. We administered a short survey to middle managers the results of which confirm the notion that supervisor turnover is a salient personal cost to most middle managers, and that they use training programs as in-kind benefits to help retain supervisors who they felt were close to quitting.

Accordingly, we interpret the results by noting that the effort of replacing supervisors who have quit and bringing new supervisors up to speed drives a wedge between the firm's objective, which is to maximize line productivity (net of turnover), and that of the middle manager. If these effort costs are high, and if the effect of training on productivity is negatively correlated with the effect of training on retention (e.g., when supervisors who will gain the most in productivity from training will be most likely to quit afterward, as would be predicted by canonical labor market models of training (Becker, 1964; Acemoglu and Pischke, 1998, 1999; Acemoglu, 1997)), middle managers might indeed recommend supervisors who they know would gain little in terms of productivity from training, but for whom training would greatly increase the probability of retention.

Analysis of impacts on supervisor turnover confirms these predictions. Training had a significant negative average treatment effect on turnover, with trained supervisors more than 14% less likely to quit than control supervisors. These impacts are driven entirely by highly recommended supervisors, who exhibited a 28% reduction in quitting, with no discernible reduction for low-recommendation counterparts. Moreover, highly recommended supervisors in the control group are more likely to quit in the absence of training than are low-recommendation controls, indicating that middle managers indeed chose their recommendations in accordance with some private knowledge of which of the supervisors in their charge were potential "flight risks" and which supervisors might choose to stay as a result of receiving the training. Taken together, the results suggest that middle manager decision-making regarding the allocation of training was driven by turnover concerns. We confirm in calculations of returns to the firm that the productivity gains we estimate (which are intent to treat estimates already net of any impacts mediated by retention) are many times more valuable to the firm than retention gains, such that middle managers' retention priorities are indeed at odds with those of the firm as conveyed in discussions with top executives.

Leveraging a recent econometric approach by Dal Bó et al. (2021), we also show that middle manager recommendations leveraged private information, and that these unobserved factors negatively predict productivity gains while positively predicting retention effects. We use the structure inherent in this approach to estimate counterfactual allocation rules, finding that middle manager ratings of line supervisors' management and industrial engineering skills (reflecting the communication and production planning skills most centrally addressed in the training) contain informative assessments of deficiencies which could have been used to allocate the training effectively. This pattern of course contrasts sharply with the null returns achieved if allocation had followed middle managers' actual recommendations instead. That is, the middle managers indeed possess private information on critical determinants of the heterogeneous returns to training and, accordingly, could have allocated the training to achieve returns well above those from random assignment. However, the endogenous "misallocation" that arises from decentralization of the allocation decision to these middle managers who prioritize retention over productivity generates negligible returns and, as a result, strong evidence against investing in scale up of the program.

3.1.1 Related Literature and Contribution

This study makes two main contributions. First, we provide a novel explanation for the persistence of managerial skill gaps, demonstrating how the misallocation of training investments within organizations can generate low average productivity returns. When decision-making is decentralized – as it often is regarding employee skill investments – and when the interests of upper and middle management are sufficiently misaligned, these investments may not be targeted toward the managers who would benefit most. This is an application of a canonical theoretical insight in the literature on decentralization within organizations (see, e.g., Aghion et al. (2014, 2021); Acemoglu et al. (2007); Bloom and Van Reenen (2011)). Perhaps the clearest example in the literature to date of this sort of misalignment of incentives from decentralized personnel decision-making in firms affecting workplace performance comes from the theory and empirical work around the Peter Principle (Lazear, 2004; Fairburn and Malcomson, 2001; Benson et al., 2019). But experimental evidence on the ways in which the trade-offs of decentralization play out in real-world settings is still limited. Atkin et al. (2017), who study the effects of incentives along the organizational hierarchy to adopt of new technology; Bandiera et al. (2021), who study the effects of increased autonomy for frontline employees in a government procurement office; Rigol and Roth (2021), who study strategic withholding of recommendations among microfinance loan officers; and Deserranno et al. (2022), who study effort provision by public health workers in Sierra Leone, are notable recent counterexamples. Our contribution is differentiated from these studies in our focus on the behavior of middle managers in a large private-sector organization. Moreover, though this notion is central to the study of issues like talent hoarding and the Peter Principle in promotion in the workplace(Haegele, 2022a; Benson et al., 2019; Lazear, 2004; Fairburn and Malcomson, 2001), to our knowledge this insight has not been applied to decision-making regarding training within the firm.[2]

Second, while much has been written on the impacts of managerial quality (see, e.g.,

[2]As has been noted in the academic literature going back at least to Mincer (1962) and Becker (1964), skill investments constitute a substantial cost, which often deters organizations from training their workers, drives them to selectively train, or to have workers self-select into training (for recent related work, see, e.g., Alfonsi et al. (2020); Caicedo et al. (2022); Sandvik et al. (2021)).

Bloom et al. (2013); Bloom and Van Reenen (2007); Bloom et al. (2016); McKenzie and Woodruff (2017); Hjort et al. (2022)), little is still known about the contribution of soft skills to managerial success, and in particular, about the effects of training managers in these skills on their teams' productivity. The closest papers to ours in this regard are Gosnell et al. (2020); Adhvaryu et al. (2022d); Alan et al. (2023); Bianchi and Giorcelli (2022); Hoffman and Tadelis (2021); Lazear et al. (2015); Adhvaryu et al. (2022a); Bandiera et al. (2020); Macchiavello et al. (2020); Giorcelli (2019). This aspect of our contribution is also related to work on the impacts of soft skills training on other types of workers, e.g., front-line workers and entrepreneurs (Deming, 2022; Groh et al., 2012, 2016; Barrera-Osorio et al., 2020; Chioda et al., 2021; Dimitriadis and Koning, 2022).[3] Our work is novel in its focus on the soft skills of managers, and on the measurement of productivity effects, which are often difficult to capture in objective terms (in addition to effects on manager/worker retention, absenteeism, compensation, and other relevant outcomes).

3.2 Context, Program, and Experimental Design

3.2.1 Context

We partnered with the largest contract manufacturer of readymade garments in India (among the top five largest garment exporters in the world), Shahi Exports, Pvt. Ltd., to implement and evaluate a management training program among production line supervisors. Given the continued labor-intensive production technology in the garment industry despite adoption of modern production concepts such as specialization, assembly lines, and lean production, garment manufacturing provides an excellent setting in which to study the impacts of training in personnel management practices on productivity.

The firm owns and operates roughly 70 factories, which produce orders for hundreds of international brands each year; our firm partner is the largest ready-made garment firm in India. There are three stages in the garment production process. First, fabric is cut and organized into bundles of sub-segments for different parts of the garment (e.g., sleeve, front placket, collar) by cutting teams. These bundles of materials are then transferred to sewing lines in which machine operators construct each portion of the garment and attach these portions together to make complete garments. Finally, the sewn garments go through fin-

[3] We are fortunate to be able to run this experiment in a setting in which we know soft skills training can be implemented effectively (Adhvaryu et al., 2022b) and the same setting from which we identified potentially productive skills to draw upon when designing the training (Adhvaryu et al., 2022d,a). Though experimental evidence of the impacts of soft skills training for managers on team productivity is a novel contribution in itself, this consistent setting allows the study of the allocation decision to be unfettered by concerns about implementation quality of the training and/or the appropriateness of the training content.

ishing (e.g., washing, trimming, final quality checking) and packing for shipment in advance of a contracted delivery date.

Across the cutting, sewing, and finishing departments representing these three stages of production, each factory employs thousands of workers allocated across tens of teams, each with at least 1 supervisor, and often several assistants of various designations (e.g., assistant supervisors, feeder, floater, captain). In smaller factories, a single cutting team and finishing team will service most or all sewing lines, but in larger factories each sewing line may have its own matched cutting and finishing teams. Each sewing line produces a unique order or style until completion, before progressing to the next contracted order.[4]

As discussed below, we randomized access to the training across supervisors from all three departments (as well as occasionally some additional supervisors deemed eligible by the factory from other support departments such as HR), but when studying impacts on productivity we focus only on the supervisors who are mapped to specific production lines for which we measure productivity (i.e., primarily sewing department in most factories, but some cutting and finishing in some factories when those supervisors are linked to specific sewing lines). Supervisors of sewing lines are assigned permanently to their line and are responsible for several key oversight tasks. First, when a new order is assigned to a line, the line supervisor must determine how to organize the production process, taking into account both the machines and workers available as well as the specific operations required and overall complexity of the garment style. This initial line architecture (known as "batch setting") is always set at the start of a new order and is rarely and minimally changed for the life of that order to avoid downtime.

Over the order's production run (lasting usually weeks, but sometimes months), if productivity imbalances or bottlenecks arise (often due to idiosyncratic worker absenteeism and/or worker-task-specific productivity shocks), sewing line supervisors will most often switch the task allocations of some set of workers across machines, or add a helper or second machine to some critical operations (often borrowing from other lines), but preserving the line architecture otherwise (Adhvaryu et al., 2022a). This recalibration of the worker-machine match (known as "line balancing"), which depends crucially on effective communication with workers and substantial monitoring effort represents one pathway by which managerial quality contributes to the marked increases in productivity seen over the life of an order in this setting (Adhvaryu et al., 2022d). The initial "batch setting" depends crucially on the planning

[4]Orders from brands are allocated across factories by the marketing department of each production division (Knits, Mens, and Ladies) based on capacity and regulatory and/or compliance clearance (i.e., whether a particular factory been approved for production for that brand given its corporate and governmental standards), and within factory, by first availability (i.e., whichever line is closest to finishing its current order when an incoming order is processed will be allocated the new order).

and organizational skills of the supervisor and contributes to vast dispersion in the productivities achieved under different line supervisors, even after accounting for garment style and worker skill and quality (Adhvaryu et al., 2020).

The managerial hierarchy of the firm involves several layers. Supervisors of production lines or teams as discussed represent the frontline of management. They report to production floor managers in larger factories with multiple floors and/or many lines on a floor, or to factory level production managers directly in smaller factories. These are the middle managers from whom we solicited training allocation recommendations. The factory level production manager works alongside the general manager of the factory who also oversees broader operations at the factory level. As mentioned above, there are roughly 60 factories with this structure, organized into 3 divisions of the firm (Knits, Mens, and Ladies) with roughly 20 factories in each division. The production and general managers of each factory report to the COO and CEO of their division. These 3 division CEOs and 3 division COOs report to the board (on which they also serve), and the board is overseen by a Managing Director (i.e., the head of the organization). Accordingly, we think of the production floor managers or factory level production managers whom we elicit recommendations from as the "middle managers" who report to division and firm level top management and who possess knowledge regarding factory level operations and work closely (daily) with the frontline supervisors among whom the training is being allocated.

3.2.1.1 External Validity

Given Shahi Exports' size and dominance in the export market, it is important to evaluate the extent to which the estimated impacts of training and underlying mechanisms are externally valid. While Shahi is indeed large relative to Indian competitors, it does have close peers across South and East Asia. Moreover, importantly for the focal subjects of this study, each Shahi factory operates like its own firm. Factory general managers have authority to set operational and managerial policy, leading to substantial heterogeneity across units. This is in part because factories are often acquired by Shahi from competitors, and existing workers, managers, and policies are left in place after acquisition. Moreover, while Shahi's *product market* power may be large, it competes in *local labor markets* with many other suppliers. Factory management told us that units often even compete for workers and managers against each other, and "poach" employees from one another as they would from any other external firm's factories. Relatedly, wages are tightly benchmarked to state minimum wages and thus similar to other suppliers. Productivity levels are also similar to other export-oriented suppliers in India and globally – average productivity for garment exporters in low-income countries is 50-55%, in line with Shahi's figures. Last, it is worth mentioning that while we are of course very sympathetic to the importance of external validity, we also recognize the trade-off that personnel economics studies often face, of having to hold certain organizational features constant in order to best study a particular economic phenomenon within the firm. Here, the fact that we know that the training is appropriate for this context, is delivered by the same team across units, and that hierarchy and responsibilities of each managerial layer are common, ensures that these features do not interfere with the study of delegation, while, as mentioned previously, still leaving many operational / managerial policies and performance to vary naturally across units. We believe our context offers a middle ground in this trade-off.

3.2.1.2 Middle Managers and Allocation Decisions

One central aspect of the middle manager's job is to nominate supervisors under their charge for a variety of investments or means of recognition as dictated by the firm's upper management. These include, e.g., awards for superior performance or commitment to the firm; promotion opportunities or additional workplace responsibilities; and, perhaps most importantly, access to training programs. Training is a key input garment firms provide with the goal of up-skilling workers and managers alike. Shahi Exports, like many manufacturing firms that are part of global supply chains, provides not only technical training but also a variety of training in "life skills" as well, usually prompted by buyers' sustainability initiatives. Adhvaryu et al. (2022b), for example, evaluates the impacts of Gap, Inc.'s P.A.C.E. program for female garment workers. The firm in general does not have the capacity to train all workers in such programs. Instead, upper management relies on middle managers to nominate program participants. Our program followed this mold and in doing so kept to the standard workplace practice of allowing middle managers to allocate training investment.

More generally, this practice is a reflection of the broader way in which most organizations deputize the allocation of training to middle management. Training is a costly resource that firms deploy strategically to boost productivity and profits. Much has been written about the potential for under-provision of training in equilibrium, especially when labor markets are frictionless (see, e.g., Becker (1964); Mincer (1962); Acemoglu (1997); Acemoglu and Pischke (1999), for some classic work in this area). Consistent with this, there is also a large ongoing policy debate globally about how to better incentivize firms to train new workers to create greater pools of skill in the economy, given firms' general sensitivities to the high costs of training (see, e.g., Alfonsi et al. (2020); Caicedo et al. (2022)). This constrained allocation of training shows up in many other contexts as well, e.g., in auto manufacturing (Adhvaryu et al., 2023) and the quick-service restaurant industry (Adhvaryu et al., 2022c).

3.2.1.3 Middle Managers and Supervisor Turnover

Given the central role supervisor retention plays in our upcoming analysis, we administered a survey in September 2021 to 50 middle managers and upper managers in 5 factories in order to better understand the roles of the middle managers in relation to supervisor turnover.[5] While the sample is not representative of our study population, it provides suggestive evidence that middle managers bear personal costs when line supervisors leave. Specifically, middle managers are personally involved in many facets of replacing and onboarding line supervisors. 70% of respondents indicate that middle managers fill in for a departed supervisor before a new supervisor is assigned. 88% indicate middle managers are involved in the replacement of the line supervisors, where involvement is broadly defined as finding, interviewing, or screening candidates. 88% also indicate that the middle managers are involved in the training of new supervisors. New supervisors do not learn the necessary skills immediately. 72% respond this process takes one week, while the remaining responded “two or three weeks” (24%) or “a month or more” (4%). Finally, all respondents reported that, from a menu of options, they would provide in kind benefits (such as training or development) programs to retain talented supervisors. To be clear, we do not argue that middle managers are the only employees involved in these processes. Survey results show that HR and upper management are also involved in replacing and training supervisors.[6] Similarly, 34% indicates floaters or assistant supervisors can be involved in filling in for a supervisor. However, collectively, the above patterns indicate that middle managers are generally involved in each step of filling in for, replacing, and training a supervisor, indicating that there is a personal cost to middle managers when line supervisors quit.

3.2.2 Program Details and Content

Drawing from our prior work in this specific context showing the productive value of both soft skills such as communication and specific managerial skills and practices such as control, autonomy, and attention (Adhvaryu et al., 2022b,d), the STITCH program was designed to train line supervisors in the skills and practices most likely to improve productivity, as line productivity has been identified by the firm as the main difference maker in profits.[7] The

[5] Of the 50 respondents, 34 are designated as floor managers. The remaining 16 are Assistant Production Managers or Production Managers, who would be above floor managers in the organizational hierarchy of a large factory but would serve as floor managers themselves in smaller factories.

[6] We also ask the respondents to rank the relative involvement of the titles they indicate as involved in a process. HR tends to be ranked lower than middle managers with regards to replacement and training, while upper management tends to be ranked higher.

[7] In discussions with senior executives of the firm, initiatives to reduce energy consumption and worker turnover are also mentioned but are considered at least an order of magnitude less impactful for profits than labor productivity.

program consisted of 25 weekly hour-long sessions divided to 4 main modules, each of which focusing on a different aspects soft-skills and leadership training. Figure C.4 presents a diagram with all 4 modules and the topics of the 25 sessions they cover. Below we give examples of how the STITCH training relates to the skills previously identified as productivity enhancing in this specific context in two studies, leaving a full discussion of the contents of each session to Appendix Section C.1.1.

First, Adhvaryu et al. (2022b) finds evidence that soft-skills training primarily focusing on effective communication, time/stress management, and problem solving makes garment sewing workers more productive, primarily through improved teamwork and collaboration. These exact skill sets are emphasized throughout STITCH training as they have the potential to be productivity enhancing with supervisors having a large role in enabling collaboration in our context. For example, in two sessions that directly emphasize communication skills, trainees participate in role-playing activities to learn about different communication styles, importance of communicating assertively and responsibly, and practice skills of giving and eliciting constructive feedback. For stress management, two sessions have explicitly focused on activities to understand emotional responses to situations, what positive actions can help manage emotions, identifying causes and effects of stress, and tips for effective stress management. For problem solving, there is a session where participants are trained on problem solving skills such as problem identification, analyzing the root cause, and making decisions from available options using case studies and role-playing exercises.

Second, Adhvaryu et al. (2022d) show supervisors with higher managerial control, attention, and autonomy enable higher team productivity. *Control* refers to the belief in the capacity to influence and control events and outcomes, which has been underlined throughout the STITCH training. Broadly, the program focuses on supervisory behaviors, skills, and attitudes that enables effective team performance and instills to the trainees that their actions strongly influence workplace outcomes and productivity.

Attention broadly refers to undertaking practices that demonstrate effort and attention to accomplishing managerial tasks. One aspect of attention is active personnel management, which is related to a session focused on effective methods of employee motivation through, for example, showing appreciation and helping employees realize their value. Frequency of monitoring work is also a key component of managerial attention, and the importance of monitoring is underlined in multiple STITCH sessions on the role of the supervisor, effective planning/organizing, and building a good work culture.

Autonomy encompasses behaviors and practices that capture the degree to which the supervisor directs the team's activities proactively and without relying on input from superiors.

The ability to do so also relies heavily on the nature of the rapport established with workers.[8] While communication related sessions discussed above have direct bearing on having a good rapport with subordinates (and to directing the team's activities), STITCH training had additional sessions focusing on the importance of sensitivity in interpersonal relationships and prevention of harassment by supervisors. With regards to directing team activities, a planning and organizing session focused on the importance of planning for effective team work and asks groups of participants to come up with a plan for their team to fulfill hypothetical orders.[9]

3.2.3 Experimental Design

Training participants were chosen from a pool of supervisors indicated by management to be eligible for training. All eligible supervisors were administered the baseline survey and were randomized into treatment and control. This gives us a baseline sample of 1849 supervisors. Employees that oversaw supervisory roles yet were not officially designated as supervisors (such as assistant supervisors or floaters) could also be indicated by management as being eligible for training. We do not make a distinction based on official designation and refer to each eligible employee as supervisors for the rest of the text. The middle managers, again as indicated by management, were also administered a baseline. In the baseline, the middle managers were asked to rank the supervisors they managed (from 1 to 5) according to how much they believed the supervisor would gain from training. The wording of the question (presented in the next section) made clear that this ranking would indeed affect the probability that the supervisors were included in the first batch of training.[10] We refer to this variable as the middle manager recommendation for the rest of the text. We collected this middle manager recommendations for 1175 supervisors included in the analysis.

Randomization was stratified in multiple dimensions. For supervisors with middle manager recommendations, whether middle manager recommendation was high, moderate, or low were used as strata. Second, for supervisors that were mapped to specific production floors, the production floor was used as a strata. Finally, supervisors were grouped into similarity clusters based on personal and line characteristics for randomization.

While randomization was at the supervisor level, our key outcome of productivity is at

[8] These two leadership styles map to the two types of behaviors identified in the leadership literature as "initiating structure" and "consideration" (Stogdill and Coons, 1957).

[9] It is important to reiterate here that though prior work from Adhvaryu et al. in this context provides a solid foundation from which the content of the training was created, this study provides the first experimental evidence on the impact of soft skills training for managers on the productivity of their teams.

[10] In practice, treatment assignment probabilities varied only slightly to be one percent higher for the highest recommendation and one percent lower for lowest recommendation supervisors, relative to supervisors with average recommendation.

the line level. Of the 1849 supervisors administered a baseline, a subset of 954 supervisors who (1) undertook duties directly related to production[11] and (2) could be linked to specific production lines, are included in our productivity analysis. These 954 supervisors were linked to 561 production lines. The line level treatment is defined as the proportion of supervisors in a line who were treated. This leads to a continuous line level treatment between 0 and 1. Treatment is evenly centered around 0.5 as shown in Appendix Figure C.3.

Appendix Figure C.1 presents a schematic diagram of our experimental design. Of the baseline sample 1849 supervisors, 921 were randomized into control and 928 were in the treatment group. Of the 561 production lines, 164 had no supervisors treated, while 397 had at least one treated supervisor. Summary statistics and balance checks are presented in section 3.4. Appendix Figure C.2 presents a timeline of our intervention. Middle managers and supervisors were administered the baseline surveys from December 2016 to March 2017. Training start and end dates were different across different factories.[12] The earliest training started in April 2017, and the latest completion was in March 2018. We discuss the various survey instruments and other data sources in Section 3.3.

This experiment predates our research team's adoption of a policy of pre-registering analysis plans for RCTs. However, we believe the structure of randomization conveys clearly at least the core analysis we intended to undertake. That is, we elicited manager recommendations and randomized within the levels of these recommendations with the explicit aim of both recovering the ATE and the degree to which managers would have allocated the program effectively for the two key workplace outcomes of interest in this setting: productivity and retention. This exercise is only relevant under a hypothesis that the treatment impacts are heterogeneous across managers. We did not have a clear prior as to whether managers would or would not allocate according to the largest ultimate productivity gains. Given the limited empirical work analyzing this problem – as well as plausible theoretical mechanisms going in both directions – we approached these as inherently empirical questions, planning to leverage the RCT's structure and accompanying data to disentangle as much as possible.

3.3 Data

We use a combination of administrative data from the factories and survey data to evaluate the program and study its allocation. We discuss these different data sources below.

[11]This excludes supervisors who, for example, are in HR and Admin departments or are data entry operators in accessory stores.

[12]While the training start dates for factories were not randomized and were partly driven by logistical reasons, within each factory the lines are randomized into treatment and control such that we can analyze differences between treatment and control at the same time relative to treatment within each factory.

Production Data. Each production line on the sewing floors records hourly productivity data. We aggregate the hourly data to the day level for each line. The key productivity measure in our analysis is *efficiency*, defined as the daily garment quantity produced over the target quantity for the day. Efficiency accounts for the number of workers on the line and the complexity of the operations performed as the target quantities are calculated by the firm using a global garment industry standard measure called Standard Allowable Minutes for each garment type. As such, efficiency captures a standardized measure of labor productivity. Any potential treatment-induced gains in efficiency can be interpreted as an increase in productivity per worker due to better managerial input. Measurement of these productivity measures have been undertaken by the firm independent of STITCH training; therefore we have access to productivity data before, during, and after the training. For productivity, our analysis period spans the 6 months before training start to 6 months after training end for each production line.

Human Resources Attendance, Pay, and Personnel Data. Human resources collects daily attendance data reporting whether an employee has attended work on a given day. We use this data to analyze supervisor attendance. We further use the attendance rosters to ascertain whether a worker is retained by the firm on a given day to investigate retention results.[13] We also have have access to monthly salary data which we use to see whether trained supervisors experience differential gross salary growth. The firm also has an incentive scheme where bonus payments are made to employees based on performance. Daily data on incentive payments are collected by the firm with the production line and designation of the individual who received the payment noted in the data. We use this daily data to assess whether workers in lines with treated supervisors receive higher incentive payments. Finally, using human resources personnel rosters, we further match approximately 55000 workers to the production lines with randomized supervisors. We use the attendance rosters data for these workers to analyze both baseline values and treatment effects regarding attendance and retention for workers.

Supervisor Baseline Survey. We administered a baseline survey of supervisors eligible for training from December 2016 to March 2017. The survey covers demographics, experience and tenure, various aspects of managerial quality and style, personality characteristics, and self assessment of skills. We use these characteristics to investigate the determinants of middle manager recommendation (discussed in the next paragraph) and the heterogeneity of treatment gains. Appendix C.4.2 provides a list of survey indices we use.

[13] While employees who quit are eventually dropped from the roster, this can happen with delay. We can use the trailing absences before a worker is dropped to pin down the effective date an employee has quit.

Middle Manager Baseline Survey. We conducted a survey of middle managers, who are above the line supervisors in the firm hierarchy. We primarily use this survey for two goals. First, and most importantly, we elicit information from the middle managers about which supervisors under them should be prioritized for training based on who they think would gain the most. We refer to this as middle manager recommendation. Second, we also elicit information from the middle managers about the managerial skills, technical skills, industrial engineering skills, and the motivation to improve of the supervisors they manage. We use these skill and motivation scores in exploring the determinants of middle manager recommendations and in exploring whether the middle managers have useful knowledge about the line supervisors they manage. Specifically, we ask the following questions to measure these features, in the following order:

1. **Skill Scores**: "Imagine a ladder with 5 steps. At the lowest step is a supervisor/floater you know, who has the lowest level of [*skill of interest*]. At the highest step is a supervisor/floater you know, who has the highest level of [*skill of interest*]. On which step would you place each of the supervisors/floaters? "

   For the following [*skills of interest*]

   (a) "technical tailoring skills"
   (b) "industrial engineering (IE) skill (e.g., assigning workers to operations, meeting targets, relieving bottlenecks, line balancing)"
   (c) "non-technical management skill (e.g., communication, leadership, ability to motivate line, sense of responsibility)"

2. **Motivation to Improve**: "Imagine a ladder with 5 steps. At the lowest step is a supervisor/floater you know, who has the lowest level of motivation to improve his/her skill. At the highest step is a supervisor/floater you know, who has the highest level of motivation to improve his/her skill. On which step would you place each of the supervisors/floaters? "

3. **Middle Manager Recommendation**: "HR is planning to train supervisors and floaters in soft skills (e.g., communication, leadership, time management, problem-solving and decision-making). Shahi Management feels it cannot train all supervisors at once and would like to focus first on those who will benefit the most. We would like to know who you think will gain the most from this training. Taking into account current skill levels and ability for improvement please rank all your supervisors in order

of who you think will benefit the most (Rank 1 is most benefited). Those who you say will have the highest expected gain will have a higher chance of getting this training."

In the rest of the paper, for consistency of exposition with other skill scores, we flip the coding of the middle manager recommendation rankings so that higher values signify a higher recommendation.[14]

Middle and Upper Manager Follow Up Survey. We administered a follow-up survey to a group of 50 middle managers and upper managers in 5 factories in September 2021. This short survey was administered to help interpret the results that highly recommended supervisors gained the least in terms of productivity and most in terms of retention (discussed in sections 3.4.2.1 and 3.4.3.1 below). It focused on the role and responsibilities of middle managers in relation to supervisor turnover. We discuss the results of this survey above in section 3.2.1.3.

Pre- and Post-Module Test Scores. Before and after each training module, all treated supervisors and a randomly selected group of control supervisors were given a short test covering the material of the module. We use the percentage point scores of these tests to assess whether treated supervisors learn the content covered in the training.

3.3.1 Summary Statistics and Baseline Balance

Table 3.1 presents summary statistics and confirms balance across characteristics of interest at the supervisor level. Given our key outcome of productivity is at the line level, we also summary statistics for production lines and confirm balance at the line level in Table 3.2. In Appendix Table C.3, we present further summary statistics and balance checks for several analysis subsets. Specifically, as we further discuss below, we drop lines with above a certain cutoff of zero productivity days in the data from our analysis in our preferred specification, as in our context this is likely a data entry error as opposed to actual zero productivity. Further, in our heterogeneity analysis with regards to middle manager recommendation, we limit our sample to lines for which we have middle manager recommendations (i.e., lines who are mapped to supervisors with middle manager recommendations). Some incidental imbalance is introduced for the subsets.[15] However, as we discuss later, we are using a

[14]The firm regularly runs surveys and training programs as a part of their standard operations, and the surveys were conducted by a firm that had been contracted with frequently by the firm. Further, the *survey* does not mention that the training is designed and evaluated by academic researchers. Therefore, we do not expect the recommendations would be different than what they would be absent the experiment.

[15]For the analysis subsample baseline productivity is 6% lower for lines with all treated supervisors (significant at 5%). For the middle manager subsample, baseline attendance is 11 % lower for fully treated lines (significant at 10%).

difference-in-differences specification for line-level outcomes, for which level differences do not pose an identification concern in the presence of parallel trends. In Appendix Figure C.7 we present event study specification results showing no evidence of pre-trends for our main analysis lines. Regardless we present results using the full set of lines available to us alongside our preferred subset to show that the coefficients are stable across samples.

Table 3.1: Supervisor Level Descriptive Statistics and Balance

	Control			**Treated**			
	N	Mean	SD	N	Mean	SD	Mean Diff. (SE)
Supervisor Age	921	31.38	6.38	928	31.16	6.03	-0.216
							(0.289)
Supervisor Male	921	0.75	0.43	928	0.74	0.44	-0.012
							(0.020)
Supervisor Finished Highschool	921	0.13	0.33	928	0.11	0.31	-0.023
							(0.015)
Supervisor Worked Different Line Before	921	0.43	0.50	928	0.40	0.49	-0.035
							(0.023)
Supervisor Ever Operator	921	0.74	0.44	928	0.72	0.45	-0.022
							(0.021)
Supervisor Worked Different Factory	921	0.26	0.44	928	0.23	0.42	-0.026
							(0.020)
Months as Supervisor	921	62.70	46.80	928	60.21	43.43	-2.495
							(2.100)
Months Supervising Current Line	920	28.87	29.09	928	29.50	30.56	0.629
							(1.388)
Years in Shahi	921	6.78	4.86	928	6.67	4.61	-0.109
							(0.220)
FIC Recommendation	591	3.08	1.62	584	3.01	1.60	-0.066
							(0.094)
Technical Skill (scored by FIC)	591	3.96	0.94	584	3.93	0.95	-0.023
							(0.055)
Industrial Engineering Skill (scored by FIC)	591	3.96	0.93	584	3.96	0.95	-0.004
							(0.055)
Management Skills (scored by FIC)	590	4.01	0.94	581	4.00	0.92	-0.010
							(0.054)
Supervisor Motivation (scored by FIC)	590	4.16	0.89	581	4.18	0.83	0.011
							(0.050)

Note: Robust standard errors are reported for difference in means (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$).

Table 3.2: Line Level Descriptive Statistics and Balance

	Num Lines	Mean	SD	Coefficient (SE)
Line Efficiency (Baseline)	540	54.96	9.31	0.940
				(0.978)
Line Attendance (Baseline)	541	0.89	0.05	-0.007
				(0.006)
Line Retention (Baseline)	528	0.83	0.14	0.004
				(0.016)
Line Budgeted Efficiency (Baseline)	541	60.60	7.58	-0.217
				(0.836)

Note: Summary statistics are included for all production lines. The coefficient(SE) is from regressing the outcome on the continuous treatment indicator. Robust standard errors are reported (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$). All baseline values are from 3 months preceding training start (January - March 2017). Baseline (budgeted) efficiency is an average of daily (budgeted) efficiency values for this period. Baseline attendance and retention are the attendance and retention outcomes for the workers matched to these lines using the personnel rosters.

3.3.2 Main Outcomes by Middle Manager Recommendation

Figure 3.1 presents the variation across high and low middle manager recommendation scores for our two primary outcome variables: line productivity and supervisor retention.[16] Panel (a) plots a histogram of baseline line efficiency for the two groups. First, the figure visualizes the substantial variation in productivity across lines. Second, the distributions largely overlap, consistent with the insignificant difference in mean baseline efficiency between the two groups (56 for low recommendation versus 55 for high recommendation lines). This implies that middle managers recommendations are not a simple function of baseline productivity of lines. We further refine this analysis in section 3.5.1.2 where we explore the determinants of middle manager recommendations.

Panel (b) plots the retention rates of control group supervisors after the start of training, again by high and low recommendation supervisors. The overall retention rate is 63%. Highly recommended supervisors in the control group are more likely to quit than their low recommendation counterparts (via hazard regression, recommended supervisors are 30% [p = 0.02] more likely to quit). This pattern suggests that middle managers may be allocating the training partly to improve retention of supervisors they believe have a higher quitting risk, a hypothesis we return to throughout the paper.

[16] Supervisors are considered high recommendation if their recommendation is above the median. In order to go from supervisor level recommendation to the line level, we average the recommendation of every supervisor tied to the line and check if the line level average is above the median recommendation of 3.

Figure 3.1: Main Outcome Variables by Middle Manager Recommendation

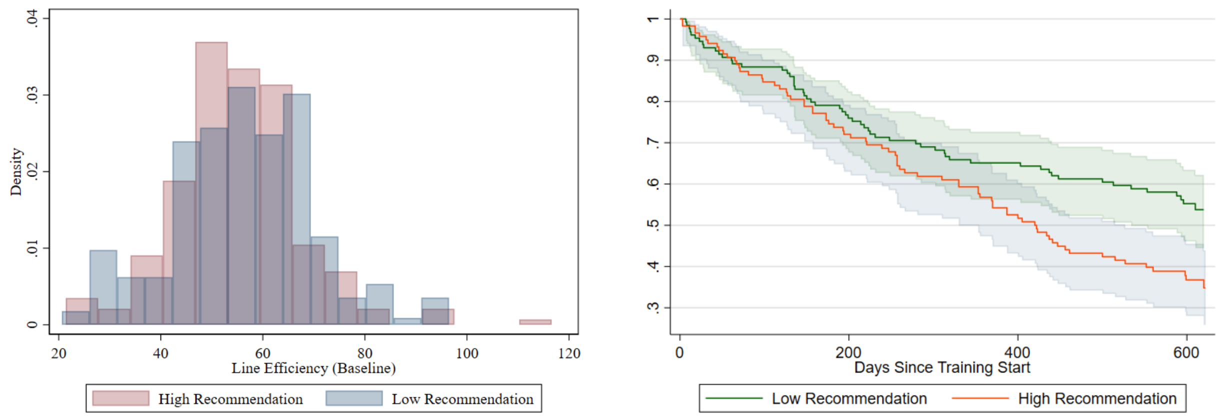

(a) Baseline Line Efficiency

(b) Control Group Retention

Note: Panel a presents baseline line efficiency values are calculated from 3 months preceding training start (January - March 2017). Panel b shows retention curves with 95% confidence intervals are shown for retention.

3.4 Treatment Effects

We start our empirical investigation by showing that the treated supervisors perform better on tests administered after each training module. Interpreting this as evidence of a first-stage effect, we next analyze the effects of the STITCH training on two key outcomes: productivity and retention. For both outcomes, we first test for average training impact. Then, we check whether supervisors recommended highly for the training by their middle manager gain more in terms of productivity/retention. We conclude the section by assessing whether these impacts are driven by spillovers within or across production lines. We leave discussion of additional treatment effects on salary growth and incentive bonus payments, important outcomes that are not central to our research design yet are inputs for our ROI calculations, to Section 3.6.

3.4.1 First Stage: Pre and Post Module Assessment

We first investigate whether the treated supervisors outperform control supervisors on tests covering the four training modules. Specifically, we compare the differences between pre- and post- module test scores for the two groups. Figure 3.2 presents the average difference for the control and treatment supervisors. While the increase in score is statistically indistinguishable from 0 across control supervisors, treated supervisors increase their test scores significantly. Appendix Table C.4 presents estimates of the treatment effect for each module using an ANCOVA specification. Consistent with the raw differences presented in the figure, treated supervisors outperform the control supervisors significantly for each module. At the low-end, treatment increases the performance of supervisors on tests of module 2 content by around 22 percentage points, which corresponds to a 40% increase from the baseline mean. The treatment more than doubles the post-training test scores for both modules 3 and 4 by inducing an increase of 32 and 39 percentage points, respectively. We interpret this as evidence that training induced substantial learning of the material covered by the modules. The impacts are not significantly heterogeneous by middle manager recommendation (Appendix Table C.5).

Figure 3.2: Pre-Post Test Score Differences

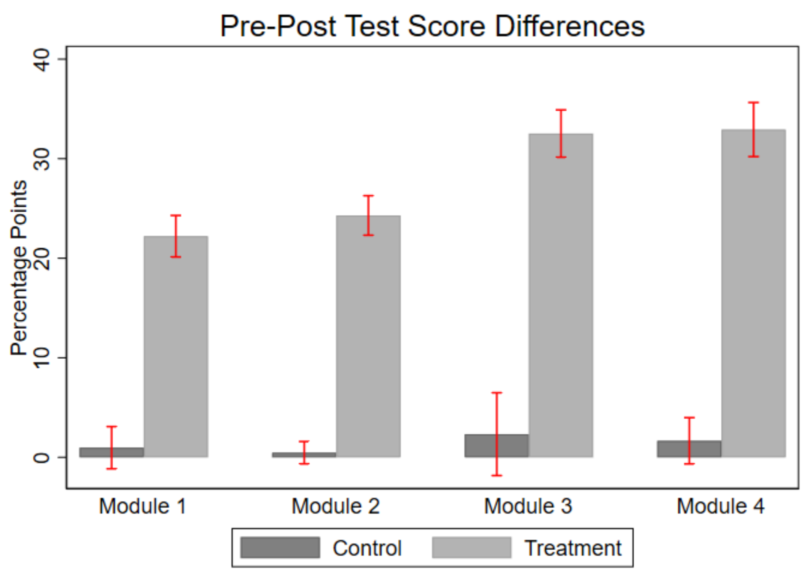

Note: Average percentage point difference between the pre- and post- module test scores for the treatment and control supervisors. 95% confidence intervals shown.

It is also worth discussing specifically how acquisition of the content of the training might have translated into the productivity gains we document below. As discussed in Section 3.2.2, the training curriculum was informed by the results of a previous study in this context (Adhvaryu et al., 2022d) which identified the skills and practices of supervisors that contributed most to line productivity. In line with these previous findings, the curriculum mainly focused on managerial *control* (belief in capacity to control outcomes and effect change), *attention* (effort and attention towards accomplishing managerial tasks, particularly personnel management) and *autonomy* (planning and organizing production proactively and without relying on inputs from superiors).

Many questions in the pre- and post-module tests gauged the internalization these particular skills and practices. For example, the test pertaining to Module 1 content asked supervisors about the value of respectfully listening to workers' perspectives, managing emotions when communicating, and effective approaches to communicating with workers including non-verbal communication. These communication styles and skills are explicit inputs into both the managerial attention and autonomy factors identified in the prior study which informed curriculum development. That is, both effective monitoring and problem identification and solving depend crucially on the quality and frequency of communication between supervisors and workers. The module 2 test measured the degree to which supervisors internalize the importance of personnel management and production planning, key components of Attention and Autonomy, respectively.

Treatment supervisors exhibited the largest gains on the tests pertaining to Modules 3 and 4. The Module 3 test focused mainly on the importance of accountability and monitoring, primary components of the managerial attention factor. While true/false questions like "changes will keep happening at work in a garment factory and managing change is the responsibility of the managers" in module 4 captured the belief in the ability to control outcomes (i.e., managerial control). In our specific context, belief in the ability to control outcomes, directing effort and attention towards personnel management, and preemptive planning and proactive adjustment to plans combine to maximize productivity. That is, Adhvaryu et al. (2022d) find that the most productive supervisors engage in close and frequent monitoring of the line, open communication with workers to understand issues and bottlenecks, balancing of lines with new workers or the reallocation of workers across operations if needed, and proactively solving issues raised by workers such as poorly calibrated machines.

3.4.2 Productivity

Next, we investigate the impact of training on line productivity. Our outcome of interest is efficiency, which is the industry standard measure of productivity defined as quantity produced over target quantity. We use the following intent-to-treat (ITT) difference-in-differences (DD) specification to assess the productivity effects on a line-day level:

$$y_{ltr} = \alpha + \beta_1 \, T_l \, \mathbb{1}[During]_{lt} + \beta_2 \, T_l \, \mathbb{1}[Post]_{lt} + \delta_l + \mu_t + \gamma_r + \epsilon_{ltr} \tag{3.1}$$

where y_{ltr} is productive efficiency of line l on date t and the relative time indicator r, T_l is treatment as defined by fraction of supervisors treated, and $\mathbb{1}[During]_{lt}$ and $\mathbb{1}[Post]_{lt}$ are indicators for whether training is ongoing or over in the factory of the line. The training-relative date r is set to 0 on the first day of the month of training start in each factory and it captures how many days have elapsed since the beginning of training in the factory of the line (with negative dates for before training). Line fixed effects δ_l controls for any time invariant line level characteristics, including the baseline characteristics of line supervisors. In our analysis sample, we exclude line-day observations with 0 efficiency as these reflect data errors as opposed to days where lines actually produced no output. Further, certain lines have many days reported as 0 productivity. We exclude any line that has over 20% of recorded days with 0 productivity in any period (pre-, during- or post- training) from the analysis. Results are stable under other reasonable cutoff values and we further show robustness to not dropping any lines or zero-productivity days in the main results Table 3.3. Our coefficients of interest are β_1 and β_2 which estimate the causal effect of fraction of supervisors treated on line level productive efficiency.

An important feature of our design is the inclusion of training-relative date fixed effects γ_r in our preferred specification. Their inclusion underlines an important distinction between our setting and the standard staggered adoption difference-in-differences or event study designs about which there has been an active recent literature (Bilinski et al., 2022). Namely, we have within cohort randomization. While different factories start treatment on different dates, within each factory lines are randomized into treatment and control. Therefore, even for $r > 0$ we have both treated and untreated production lines. Inclusion of training-relative date fixed effects γ_r allows us to recover treatment effects by comparing treatment and control lines within cohort and training-relative date. The recent literature has raised concerns about the interpretation of two-way fixed effects estimators in the traditional staggered adoption setup (de Chaisemartin and D'Haultfœuille, 2020; Borusyak et al., 2022c; Goodman-Bacon, 2021; Sun and Abraham, 2021b) in the presence of treatment effect heterogeneity across either time or treatment cohorts or both. Within factory randomization of production lines into treatment and the ability to recover treatment effects within cohort-time alleviates these concerns. We further show the robustness of our DD productivity results to running the analysis on a balanced panel in relative time (6 months before to 20 months after training start for each line) in our robustness discussion below.

As shown in Table 3.3, treatment has a statistically and economically significant effect on efficiency. Column 3 reports our preferred specification, which includes line, date, and relative date fixed effects (columns 1 and 2 show robustness to including a less stringent set of fixed effects). Lines with all supervisors treated are, on average, are 4.1 percentage points (7.3% of control mean) more efficient during training. Given the training was administered over a considerable duration of an average of 9 months, this is an economically significant effect. For the 6 months following training end, lines with all supervisors treated still have

3.3 percentage points (5.8% of control mean) higher efficiency then lines with no supervisors treated. This implies that while the productivity impact of the training is stronger during the lengthy training period, the effects persist after training completion. The decreasing line-level productivity effects over time makes sense given we train individual supervisors, yet analyze productivity at the level of the supervisors' line during randomization. As supervisors can leave the firm or possibly get reassigned to different lines (or successful practices can spillover to control lines) over time, we might expect the line-level treatment effect from treating supervisors dampen over time. However, the results show that the sum of these possibilities is not enough to erode the treatment effect substantially for at least 6 months after training completion. The average productivity results strongly indicate that the training increased profitability of the firm, as the firm focuses on labor productivity as the main lever for affecting profits. That is, the other major sources of costs are raw inputs such as cloth, yarn, and energy over which the firm feels it has little control. We confirm this intuition in section 3.6 by undertaking a returns on investment analysis using information from the firm's accounting department.

Table 3.3: Effects of Training on Line Productivity

	Dependent Variable: Efficiency (Produced/Target)					
	Analysis Lines			Lines w/ Middle Manager Match	All Lines	All Sup. Treated or Control
	(1)	(2)	(3)	(4)	(5)	(6)
During Training X Treatment	3.986***	3.994***	4.089***	3.873***	4.208***	3.967***
	(1.113)	(1.115)	(1.116)	(1.249)	(1.286)	(1.354)
After Training X Treatment	3.079**	3.075**	3.267**	3.258**	3.296**	3.597**
	(1.364)	(1.366)	(1.338)	(1.564)	(1.632)	(1.683)
Observations	228167	228166	228166	189380	254138	151104
Number of Lines	480	480	480	395	553	314
Cont. Mean of Dep. Var.	55.865	55.865	55.865	55.865	55.865	55.462
Line FE	X	X	X	X	X	X
Month FE	X					
Day FE		X	X	X	X	X
Relative Date FE			X	X	X	X

Note: Standard errors are clustered at line level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). The analysis covers six months prior to training start month and the six months post the training end month for each factory. For columns (1) - (4) days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis. Column (5) includes both the dropped lines and the production days with 0 efficiency. Column (6) only includes the analysis lines where all supervisors are either treated or control.

To assess the dynamics of the productivity results, the Appendix Figure C.7 presents the monthly event study results, starting from 6 months before training start to 20 months after training start for each line. The event study shows no clear pre-trend that should cause concern for identification. It further provides hints about the dynamics of the treatment effect. The treatment effect rises the first 4 months after training start and peaks at around month 4 on average. After that, the treatment effects start getting smaller, but coefficients stay positive for the rest of the analysis period. This suggests that, even after close to a year after training end for many lines, the treatment effects do not go to zero. The fall in treatment effect size (about 30%) in this time period is consistent in magnitude with the attrition of treated supervisors 15 months after the end of training (about 35%).

Robustness. Column 4 of Table 3.3 shows that our results hold in the subset of lines with middle manager recommendation information. Column 5 shows robustness to including both the 0 productivity line-day observations and the lines with many 0 productivity days in our analysis. Column 6 further assuages concerns regarding bias due to within-line variation in treatment, by showing that our results are unchanged if we focus on the lines where all supervisors are either treated or control. While results in Table 3.3 are from an unbalanced panel in relative time due to differing treatment length across factories, Appendix Table C.6 shows that the results are robust to running the analysis on a balanced panel in relative time (6 months before to 20 months after training start for each line). While there is an approximately 10% decrease in the after-training effect size, this is to be expected as the analysis period includes more months farther away from training for every line. Finally, Appendix Section C.5.5 shows robustness of our productivity results to defining line-level treatment as treating any supervisor on a line.

3.4.2.1 Productivity Effect Heterogeneity by Middle Manager Recommendation

Central to our research question is whether the middle managers would allocate the training to supervisors who would gain the most. If middle managers possess private information about who would gain the most from training, they could allocate training to maximize gains. However, middle managers can also have different objectives than the firm, which could conceivably lead to allocation rules that do not maximize productivity gains from training. To start exploring this question, we test if the productivity gains are higher for lines with highly recommended supervisors by their middle managers. To do so, we modify our difference-in-differences specification in equation 3.1 to include three way interactions between treatment, the treatment periods, and high middle manager recommendation indicator ($\mathbb{1}[Rec_l]$). In order to go from supervisor level middle manager recommendation to the line level, we average the recommendation of every supervisor on the line. We set $\mathbb{1}[Rec_l] = 1$ if the line level average is above the median recommendation rank of 3.[17]

Figure 3.3: Line Productivity Heterogeneity by Middle Manager Recommendation

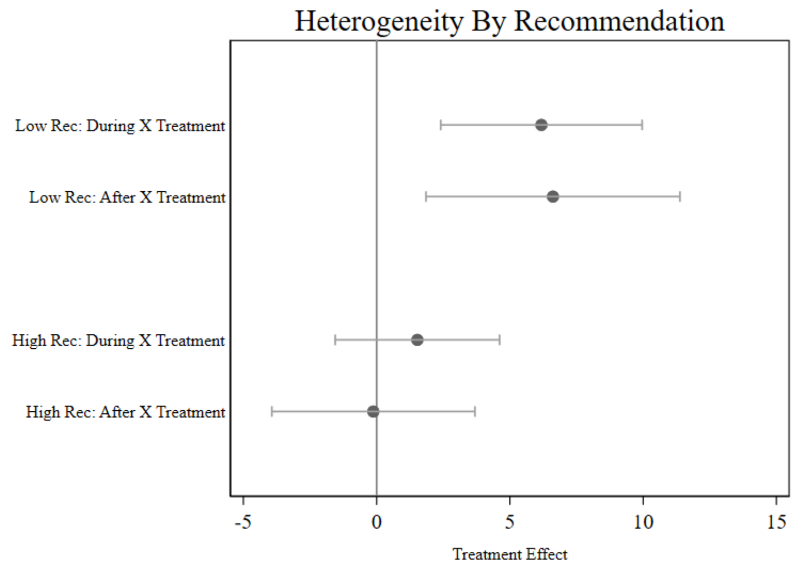

Note: Treatment effects on line productivity for high and low middle manager recommendation lines. A line is defined as high recommendation if the average recommendation of the supervisors on the line is above the median. 95 % confidence intervals are shown. Column 3 of Appendix Table C.7 reports the underlying regression results.

Figure 3.3 charts the treatment effect estimates for the high- and the low-recommendation lines. It is clear that there is significant heterogeneity in how much different supervisors gain from treatment. Strikingly, the productivity gains are highly concentrated among lines with low recommendation supervisors. During training, treated lines with low recommendation supervisors experience a 6.2 percentage points (11% relative to baseline) increase in productivity. For lines with recommended supervisors, the corresponding treatment effect is only 1.5 percentage points (3% relative to baseline) and statistically insignificant. 6 months following treatment, lines with highly recommended supervisors effectively do not gain from training while low recommendation lines have a treatment effect of 6.6 percentage points. Appendix Table C.7 reports the underlying regression results (column 3) showing the difference is statistically significant, along with robustness to less stringent FEs and using all production lines in our analysis. Finally, given production lines can have both high- and low-recommendation supervisors, we confirm that our choice of aggregation does not drive

[17]As mentioned in Section 3.3, we flip the middle manager recommendation rankings in order for higher values to mean higher recommendation.

these results by showing that the treatment heterogeneity remains unchanged when we only include lines with either all high- or low-recommendation supervisors (Appendix Table C.8, Column 6).

3.4.3 Supervisor Retention

Next, we focus on the impacts of training on supervisor retention. We estimate a Cox proportional hazard model, taking the randomization strata into account:

$$q_{ist} = h_{0t} \, e^{\mu_s + \beta T_i} \tag{3.2}$$

where q_{ist} is the hazard function for quitting, μ_s is the randomization strata fixed effects and T_i is the treatment indicator for supervisor i at time t. We present results for both the full set of supervisors in our study and for the subset of supervisors for whom we have middle manager recommendations, as this is the subset we use to assess heterogeneity below. We limit our sample to supervisors who are with the firm when the training starts in their factories. The analysis spans from the first day of training to end of 2018.

Table 3.4: Supervisor Retention

	Dependent Variable: Supervisor Quit	
	All Supervisors	Supervisors w/ Middle Manager Rec.
	(1)	(2)
Treatment	-0.151**	-0.081
	(0.073)	(0.098)
Observations	1419	889
Relative Hazard of Treatment	0.859	0.922
Strata FE	Yes	Yes

Note: *** p<0.01, ** p<0.05, * p<0.10. The sample is restricted to supervisors that could be matched to the attendance roster and supervisors who did not quit the firm between the baseline survey and the training start in their factories. Column 2 further restricts the analysis to supervisors for which we have middle manager recommendations. Relative Hazard is calculated as the exponent of the coefficient on treatment.

Retention results are presented in Table 3.4. Both in the production sample and the full sample, treatment leads to a decrease in the hazard ratio for quitting. In the full sample, treated individuals are 14% less likely to quit. For the smaller sample with middle manager recommendations, the treated supervisors are 8% less likely to quit. While economically meaningful, it is imprecisely estimated and not statistically significant, unlike the results for full sample. It also masks significant heterogeneity by middle manager recommendation, which we turn to next.

3.4.3.1 Retention Effect Heterogeneity by Middle Manager Recommendation

To explore whether the retention effects are heterogeneous with regards to middle manager recommendation, we plot survival curves (with retention as the outcome) for treated and control supervisors, separately for high- and low- recommendation supervisors. Figure 3.4 presents the results, with Figure 3.4a showing the curves for high-recommendation supervisors and Figure 3.4b include the low-recommendation supervisors. The retention effect of training is driven entirely by highly recommended supervisors. Appendix Table C.14 show results from the associated proportional hazard models, showing treated supervisors are 28% less likely to quit among high-recommendation supervisors, while there is no discernible effect for low-recommendation supervisors. Note that this pattern suggests that retention effects are not a major source of productivity gains from the training as the retention gains are concentrated among supervisors with high middle manager recommendations, precisely the group that exhibits negligible productivity gains.

Figure 3.4: Retention Treatment Effects for High and Low Middle Manager Recommendation Supervisors

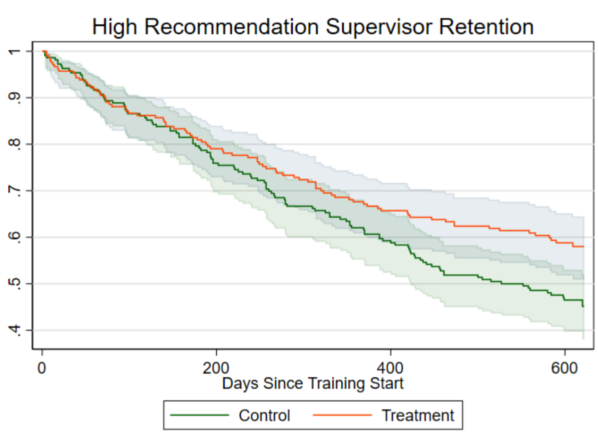

(a) High Middle Manager Recommendation

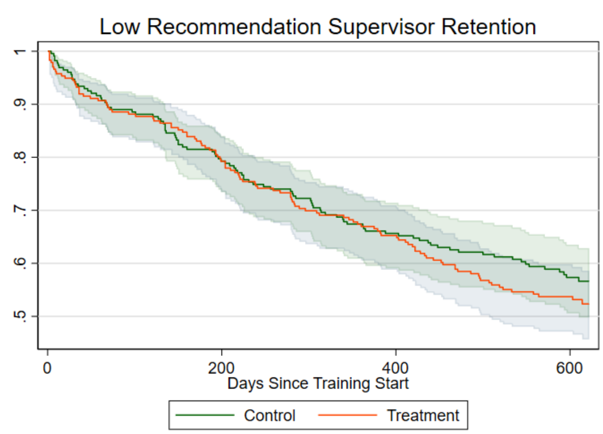

(b) Low Middle Manager Recommendation

3.4.4 Are Treatment Effects Driven By Spillovers?

Given treatment is randomized and assigned at the level of individual supervisors, our study features (1) some production lines with both treated and control supervisors, and (2) production lines with varying treatment levels co-existing in close proximity on production floors. In this section, we assess (and fail to find evidence for) whether the documented impacts can be driven by within- or across-line spillover effects.

Within-Line Spillovers. Negative spillovers within production lines can lead to biased estimates of productivity and retention gains. For example, control supervisors who are working among trained supervisors in the same line and know that they were passed over for training may be affected negatively themselves. If these supervisors quit due to demotivation or career concerns, the treated supervisors' turnover may look lower by comparison. To see if such a mechanism is at play, we compare the retention of control supervisors with and without trained supervisors in their same lines. Table 3.5 presents results from a Cox proportional hazard model that fails to find evidence that a control supervisor having any treated supervisor on the same line (Column 1) or the share of co-supervisor who are treated (Column 2) leads to higher attrition. Appendix Figure C.9 further plots survival plots, showing no differential retention pattern between control supervisors with and without treated supervisor on the same line. These findings point against a substantial negative within-line spillover to control supervisors.

Table 3.5: Control Supervisor Retention by Linemate Treatment

Dependent Variable: Supervisor Quit		
	(1)	(2)
Any Treated Linemate	0.073	
	(0.169)	
Share of Linemates Treated		-0.105
		(0.227)
Observations	358	358
Relative Hazard of Treatment	1.076	0.900
Strata FE	Yes	Yes

Note: *** p<0.01, ** p<0.05, * p<0.10. The sample is restricted to control supervisors that could be matched to the attendance roster and those who did not quit the firm between the baseline survey and the training start in their factories. Relative Hazard is calculated as the exponent of the coefficient on treatment. *Any Treated Linemate* is an indicator for having any other treated supervisor on the line. *Share of Linemates Treated* is the share of other treated supervisors on the line who are treated. For supervisors without any other supervisors on the line, both values are set to zero.

We are unable to undertake an analogous analysis for productivity because we do not have supervisor-level productivity estimates. Instead, we check if productivity effect estimates change if we restrict our sample to production lines where there is no scope for within-line spillovers. Table 3.3 Column 6 shows that the productivity effect estimates are unchanged when we restrict the analysis to the 314 production lines where all supervisors are either treated or control. The stability of the estimate between between the full and the restricted sample points against substantial bias in productivity estimates stemming from within-line spillovers. Finally, Appendix Table C.8 shows that the productivity impact heterogeneity by middle manager recommendation remains unchanged as well when we only include the same 314 production lines without scope for within-line spillovers. This assuages concerns regarding whether spillovers within production lines may explain the differences in productivity gains among high-recommendation and low-recommendation lines due to, for example, differential morale impacts on control supervisors resulting from the recommendation status of their treated co-supervisor on the line.[18]

Across-Line Spillovers. Many production lines co-exist in the production floors in close proximity and these lines may influence each other in various ways. Therefore, across-line productivity spillovers from training may be present in our context. For example, if training leads to differential worker mobility patterns across lines in a production floor due to improvements in efficiency freeing workers to be transferred to lines with high absenteeism, the effects of training may reverberate across neighboring production lines. The diffusion of successful managerial practices across the production floor can further lead to positive spillover effects. Conversely, just as mentioned above, we might be concerned that the allocation of treatment to a subset of supervisors can lead to discouragement of the control group which negatively influences their productivity, or that middle managers can direct resources and attention to production lines with trained supervisors. If present, such negative spillovers can lead us to erroneously conclude that training leads to an increase in production, while in actuality what we observe is a decrease in productivity of the control group.

Our experimental design originally meant to study these spillover effects by randomly varying production floor level treatment saturation for a subset of the lines. For this subset, the original design was to have floors with 70% of the supervisors treated (high saturation) and 30% of the supervisors treated (low saturation). This design was imperfectly implemented due to complications with the mapping of production lines to floors at the time of the randomization, leading to considerable variation in the fraction of supervisors treated

[18]This is consistent with the fact that the recommendation status of supervisors is not readily observed except to the middle manager providing the recommendation.

within saturation groups.[19] However, we still use the variation induced on the fraction of supervisors treated on a floor (saturation) to check for spillover effects. First, we divide the 54 production floors with saturation variation into three groups based on terciles of saturation level.[20] We then update our main productivity specification equation 3.1 to include the triple interaction between the training periods, floor level saturation tercile, and the line level treatment. This analysis is informative about spillovers under the reasonable case that the probability and/or magnitude of across-line spillover effects are monotonic functions of floor-level treatment saturation, i.e. control supervisors exhibit differential impact from the intervention based on how many treated supervisors they have on their floor.

Figure 3.5 presents the results of floor saturation on productivity for lines without treated supervisors. It charts the productivity effects of being on the second and third terciles of floor saturation relative to the first tercile. There is no evidence of spillovers during training and suggestive evidence for positive spillovers after training. Specifically, the coefficient for being on the second and third terciles after training are 3.7pp (statistically significant at 10%) and 3.6pp (not statistically significant), respectively. While only suggestive, these results point against negative spillovers across production lines. This largely rules out the possibility that the productivity effects we observe are primarily driven by negative spillovers to the control group, for example due to a perception of training allocation as unfair.[21]

3.4.5 Summary of Main Results

Before discussing the interpretation of the results, we quickly summarize the pattern of main impacts. Altogether, we have three key takeaways with regards to productivity and retention impacts of STITCH. First, there are large average impacts from STITCH training, especially for line-level productivity. Second, there is heterogeneity in who gains the most from training, across both productivity and retention. Third, middle managers target supervisors who gain the most in terms of retention while gaining very little in terms of productivity. We now turn to interpretation of these patterns.

[19] Some production lines were not properly mapped to production floors at the time of randomization. Supervisors of these lines were not randomized with the correct treatment probability corresponding to the saturation group of the floors on which they worked. Mapping these supervisors to the correct floors after randomization showed that balance in the randomization was preserved but substantial dispersion of floor level treatment saturation existed beyond the high and low saturation amounts intended.

[20] The first tercile covers saturation levels from 0 % to 33 %, the second tercile covers 35% to 63 %, and the third tercile covers 66 % to 100 %.

[21] Of course, spillovers could be non-monotonic, with negative spillovers in some levels of floor-level treatment intensity and positive in others. We are unfortunately under-powered to undertake this much more granular study of spillovers that would uncover spillovers at different intensities, and leave such an exercise for future work.

Figure 3.5: Productivity Effects of Treatment Saturation

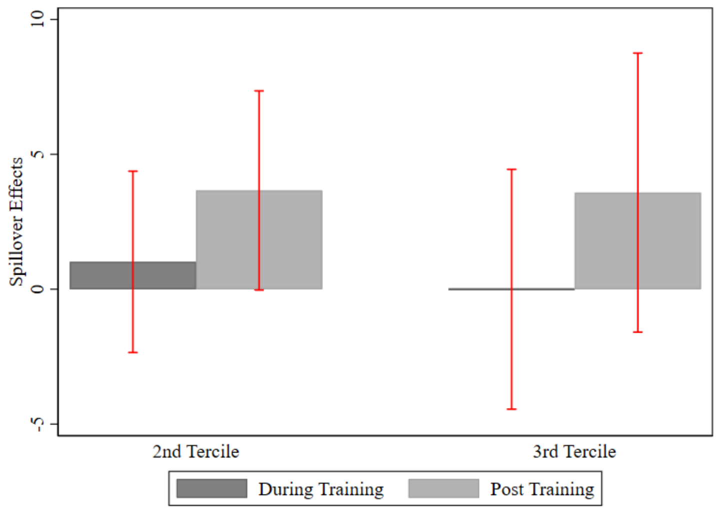

Note: Estimates underlying the figures are presented in Appendix Table C.15 column 3. The bars depict the effects of saturation on productivity for lines without treated supervisors, reporting coefficients on the interaction between during/after training and second/third saturation terciles.

3.5 Interpreting the Middle Manager Recommendations

The previous section highlights that supervisors with high middle manager recommendation (1) gain less in terms of productivity and (2) gain more in terms of retention from the STITCH training. Keeping in mind that our productivity results are ITT at the line level (such that line level productivity is observed and analyzed even if a study supervisor leaves the firm), and hence provide us with treatment effects on productivity that incorporate any effects on productivity mediated through effects on retention, we interpret these results as evidence that the middle managers were targeting retention gains above and beyond productivity gains when allocating training. This is consistent with the idea that the middle managers have private incentives to target retention beyond its implications for productivity, even if the firm would rather allocate training to increase productivity itself. In our context, there are potentially sizable personal costs to middle managers in the event of supervisor turnover. First, as survey evidence we discussed in Section 2.1.1 indicates, there are costs to middle managers in terms of time and effort to replace and onboard new supervisors. Further, middle managers can face professional costs if the line supervisors they manage quit frequently. With personal costs to retention and a trade-off between productivity and retention gains, we can rationalize the allocation decisions of the middle manager.

In Appendix C.2, we present a simple principal-agent framework that matches key elements of our context to formalize how, in the presence of private costs to supervisor turnover, middle managers could misallocate the training from the perspective of a firm that primarily targets productivity (net of turnover). We model how a firm (principal) and middle manager (the agent) would choose to allocate training to supervisors with heterogeneous gains. Training affects both productivity and retention. We allow for retention of supervisors to also indirectly affect line productivity through, for example, the replacement supervisors being less productive. The firm's objective is to allocate the training to maximize the productivity gains from training (taking the productivity effects of retention into account). However, middle managers face personal costs to supervisor turnover, creating a wedge between the firm and middle managers value of training a supervisor. If productivity and retention gains are negatively correlated – as our empirical evidence indeed suggests – middle managers would choose to allocate the training to supervisors with relatively higher retention and lower productivity gains. This would lead to misallocation from the perspective of the firm. A high personal supervisor turnover cost would lead to middle managers heavily targeting retention, leading to the observed results that middle manager-recommended supervisors gain relatively little in terms of productivity but are more likely to be retained.

The remainder of this section focuses on further exploring the middle manager recommendation choices. We start by showing that middle managers indeed possess valuable information about line supervisors they manage, implying they have private information to target the training. We then use our rich baseline characteristics data to show the observable characteristics of highly recommended line supervisors. However, much of the variation is not explained by observables, which suggests (unobservable) private information drives much of the recommendation decisions. We adapt the framework proposed by Dal Bó et al. (2021) to decompose the middle manager selection to observable and unobservable components and show that the unobservable component of selection primarily predicts lower productivity gains and higher retention gains for recommended supervisors. We conclude the section by exploring alternative explanations that can potentially rationalize our middle manager recommendation findings.

3.5.1 Who Do the Middle Managers Recommend?

3.5.1.1 Middle Managers Have Useful Information About Line Supervisors

Before focusing on the determinants of middle manager recommendations, we first explore whether middle managers seem to have useful information on the supervisors which would allow them to allocate the training effectively. If middle managers lack information entirely, we might expect that the recommendations are effectively random. However, this interpretation is hard to square with the strong relationship between who the middle managers recommend and gains across productivity and retention. Nevertheless, we check whether the skill scores we elicited from middle managers about the supervisors support the notion that middle managers have useful information about their supervisors.

In the middle manager baseline survey, we asked the middle managers to score (from 1 to 5) all the supervisors they list as reporting to them in three dimensions: management skills, industrial engineering skills, and technical skills. In Figure 3.6, we show the correlation between all three of these scores and the baseline productivity of the line(s) they supervise.[22] The results indicate that the skill scores are positively correlated with baseline productivity, with the relationship more pronounced for industrial engineering and technical skill scores. Appendix Table C.13 further shows that the association is positive for all skills and statistically significant for industrial engineering and technical skills. This is suggestive that the middle managers have useful information about the skill sets of their supervisors.

We further investigate the information content of the middle manager skill assessments by focusing on the heterogeneous productivity gains from training with regards to the three skill types. We augment our difference-in-differences specification in 3.1 to include the three way interactions between treatment, treatment periods, and the average skill scores of the line supervisors. If the skill scores capture meaningful information about the skill sets of the supervisors which are then augmented by training, we would see differential effects of training by baseline stock of skills as indicated by middle managers. Further, given the focus of the training is managerial skills, we would expect that baseline level of managerial skills and industrial engineering skills (which most closely map to the content of the training) would be be particularly related to treatment gains.[23] Table 3.6 presents our results. As expected, baseline industrial engineering and management skills have a large and significant effect on productivity gains. Specifically, supervisors with lower level of baseline skills gain more from the training, indicating that training is a substitute for baseline skills in these dimensions as assessed by the middle managers.[24] Technical skills are also negatively related to treatment gains, but the effect is smaller and statistically insignificant (especially during training). These results suggest that middle managers have nuanced information about the

[22]The productivity of the line is calculated following an AKM-style two-way fixed effect model, following Adhvaryu et al. (2022d). We implement the model on productivity data from the three months preceding survey (January-March 2017). Further description of the procedure can be found in Appendix Section C.4.1.

[23]Industrial engineering skills we underline in our middle manager survey includes ”assigning workers to operations”, ”meeting targets” and ”line balancing” which are skills that are related to managerial and leadership capacities covered in the training.

[24]This is consistent with Adhvaryu et al. (2022b)’s conclusions about a similar soft-skills training aimed at sewing workers.

Figure 3.6: Middle Manager Assessment of Skills and Line Productivity

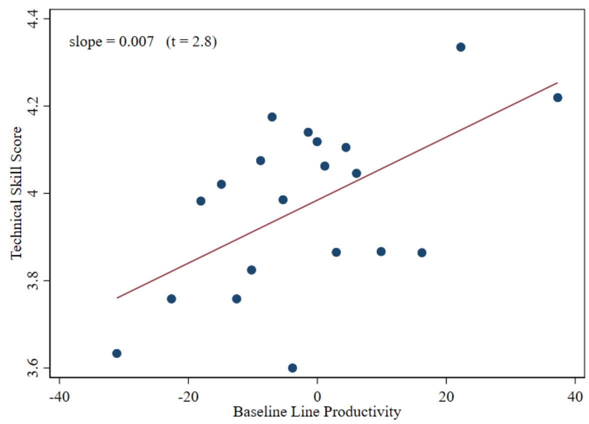

(a) Technical Skills

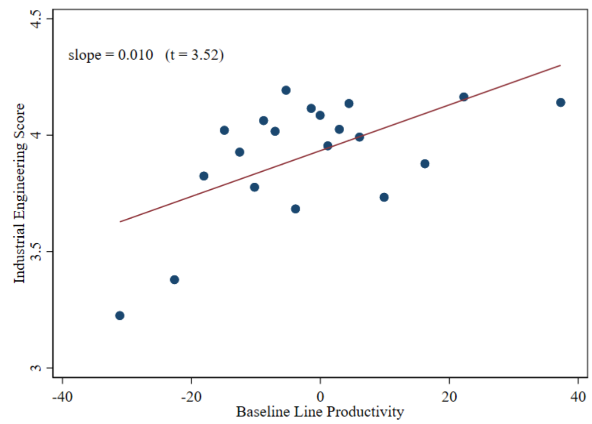

(b) Industrial Engineering Skills

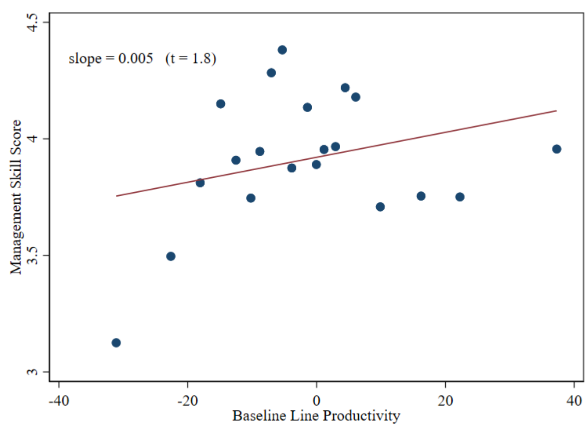

(c) Management Skills

Note: Binned scatter plots between the middle manager assessment of supervisor skill and the line productivity at baseline. Line productivity is calculated from a two-way fixed effect model matching production lines and order styles.

baseline skill sets of the supervisors they manage and this information can be leveraged to allocate training to maximize productivity gains relative to full randomization.

Finally, we remind the reader that Figure 3.1b shows that highly recommended supervisors in the control group are 30% more likely to quit than their low recommendation counterparts. That the middle manager recommendations are predictive of future quitting rates in the control group is consistent with the idea that middle managers possess private information about their supervisors. It is further consistent with the idea that middle managers could be using the training to target retention gains, as the recommended supervisors are more likely to quit the firm absent the treatment.

Table 3.6: Heterogeneous Productivity Effects by Supervisor Skill

Dep Var: Efficiency (Produced/Target)			
	(1)	(2)	(3)
During Training X Treatment	6.917	17.576***	16.716***
	(6.430)	(6.280)	(5.851)
After Training X Treatment	12.774*	18.362***	19.349***
	(6.955)	(6.678)	(6.229)
During Training X Treatment X Technical Skill	-0.818		
	(1.525)		
After Training X Treatment X Technical Skill	-2.476		
	(1.688)		
During Training X Treatment X Ind. Eng. Skill		-3.679**	
		(1.513)	
After Training X Treatment X Ind. Eng. Skill		-4.138**	
		(1.604)	
During Training X Treatment X Management Skill			-3.605**
			(1.428)
After Training X Treatment X Management Skill			-4.564***
			(1.550)
Observations	189380	189380	189380
Number of Lines	395	395	395
Control Mean of Dependent Variable	55.279	55.279	55.279

Note: Standard errors are clustered at floor level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). Line, date and relative date FEs are included. The analysis covers six months prior to training start month and the six months post the training end month for each factory. Days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis.

3.5.1.2 Observable Determinants of Middle Manager Recommendation

Having established that middle managers seem to possess information to select which supervisors would gain more from training in terms of productivity, we now turn to who the middle managers have actually recommended. Specifically, we focus on the observable determinants of middle manager recommendation. Given the rich set of baseline information we have about the supervisors and middle managers, we use a simple LASSO procedure to see which variables out of 52 supervisor characteristics and supervisor-middle manager joint characteristics (for example, whether they share the same religion or are from the same state) are associated with high recommendations.[25] Given we are primarily interested in the negative relationship between recommendations and the productivity gains from training, the analysis below focuses on the production sample (i.e. the set of supervisors who undertook duties directly related to production in a specific line) and includes line-level covariates such as baseline productivity of the line. The outcome of interest is middle manager recommendation (ranging 1 to 5).

Figure 3.7: Lasso Selected Variables

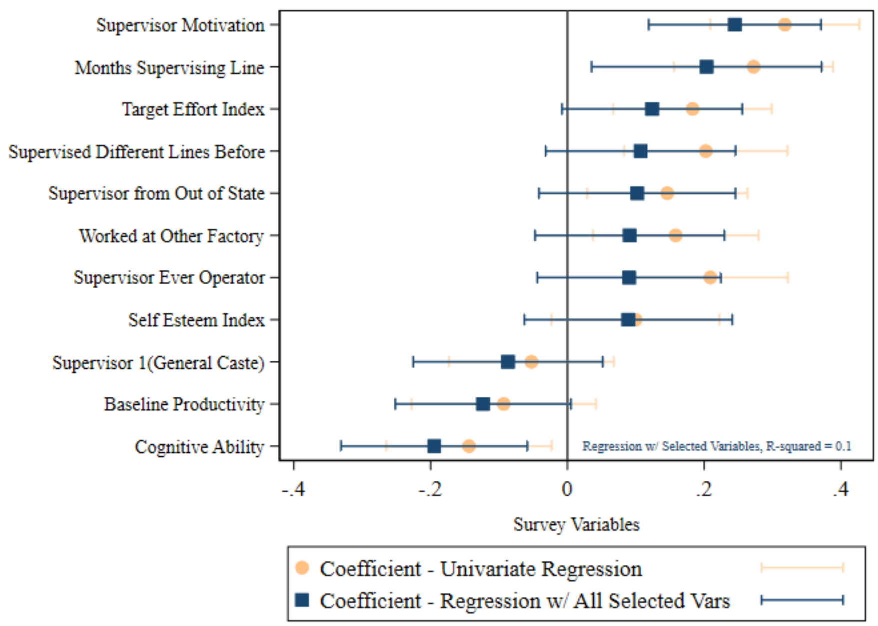

Note: Variables selected from the lasso procedure. The light orange coefficients are from a regression of middle manager recommendation on the selected variable of interest. The dark blue coefficients are from a regression of middle manager recommendation on all selected variables. 95% CI are shown using robust SEs.

Figure 3.7 presents the results. 11 variables are selected using this procedure. Note that linearly regressing the middle manager recommendations on all the selected variables yields an R^2 of only 0.1. Further, 8 out of the 11 variables have coefficients insignificant at the 5% level. We take these as evidence that, on average, much of the variation in the middle manager recommendations are driven by unobservable factors, even with the rich set of baseline variables we observe.

While our main takeaway is the overall importance of unobservable factors, a few patterns do emerge with regards to observable characteristics. First, variables indicating high

[25] The full set of included variables and additional details of the procedure can be found in Appendix Section C.4.2.

tenure and variety of experience consistently predict higher middle manager recommendation. These include months supervising current line, whether the supervisor has worked in a different line or factory before, or whether the supervisor has ever been an operator. Second, middle managers seem to recommend individuals who they view as motivated, as evidenced by not only the positive coefficient on supervisor motivation, but also that on the target effort index from the management style survey. Finally, high baseline productivity of the supervisor's line and the supervisor's cognitive ability (measured by arithmetic and digit span recall tests) predicts lower middle manager recommendations, implying supervisors may view the training as substitutes to baseline stocks of these characteristics. The fact that highly recommended supervisors have, on average, lower baseline productivity and higher retention gain implies that middle manager recommendations would increase the likelihood of retaining below-average supervisors in terms of baseline productivity.

3.5.1.3 Decomposing Middle Manager Selection

The previous section established that much of what drives the middle manager recommendation is unobservable to the econometrician, even with the rich baseline data we collected. In this section, we use a simple framework to decompose middle manager selection into observable and unobservable components and to investigate which components drive the positive/negative relationship between middle manager selection and retention/productivity. We discuss the framework and results broadly here and relegate the details of the framework and related tables and figures to Appendix Section C.3.

Middle managers who have perfect information about a supervisor's gain from training make allocation decisions based on both the gain from training as well as other manager preferences.[26] Both gains from training and manager preferences are a function of observable supervisor characteristics and a jointly normally distributed unobservable (to the econometrician) term. The key to our approach is that we allow the unobservable training gain to be negatively correlated with unobservable manager preferences. The framework leads to a two-stage estimation procedure, in which the first stage is a probit on supervisor selection (recommendation scores), and the second stage is a heterogeneous treatment effect model that leverages heterogeneity by supervisor observables and manager selection.[27] Randomiza-

[26]The framework is similar to that of Dal Bó et al. (2021). However, our setup differs from it in one key dimension. In Dal Bó et al. (2021) what leads to a possibly null relationship between the agent's selection and productivity gains is information frictions, where the agent only has imperfect information about productivity gains. If the agent's signal is very weak, her selection may not be related to treatment gains. We assume perfect information, consistent with our results that middle managers have valuable information, but instead allow for productivity-related and productivity-unrelated unobservables to be negatively correlated.

[27]For this analysis, we use a single post training-start indicator as opposed to estimating separate treatment effects for during and after training periods.

tion within middle manager recommendation bins allows identification of the second stage regression coefficients, breaking the dependence between manager selection and treatment status. We do this analysis for both retention and productivity. The second stage results suggest that the unobservable component of selection is negatively related to productivity gains and positively related to retention gains.

We use the estimated model to obtain predicted treatment effects for different allocation rules to investigate the role of the unobserved component of middle manager recommendations on the pattern of treatment effects. We present treatment effects under three alternative allocation rules following Dal Bó et al. (2021): (1) random assignment, (2) assignment based on middle manager recommendations, and (3) assignment based on middle manager recommendation with the effects of unobservables shut down.[28] The third allocation allows us to assess whether middle managers private knowledge of the supervisor unobservables is a driver of the treatment effects that would be generated under middle manager allocation. Appendix Figure C.5 presents the resulting treatment effects for line productivity and retention outcomes.

We first focus on retention. The model implies that, if half of the supervisors are treated, allocating the training based on middle manager recommendations outperforms random allocation by 68% (in terms of probability of quitting). Importantly, when we shut-down the effects of unobservables on the treatment effects we get that middle manager recommendation yields treatment effects only 28% higher than random allocation. This suggests that the private information of the middle managers allow them to target supervisors with higher retention response to treatment. The firm would presumably not be able to replicate this allocation based on observable characteristics alone (even though observed characteristics in this instance include costly measures like skill scores elicited from the middle managers).

For line productivity, as expected from earlier analysis, the model suggests that random allocation substantially outperforms allocation by middle manager recommendation. With half of the lines treated, the average treatment effect is approximately 0 with middle manager allocation, while the random allocation yields an average treatment effect of 2.8 percentage points. With the effects of unobservables shut down, middle manager allocation yields an average treatment effect of 1.6. This suggests that private information of middle managers, which allows them to target retention, is negatively correlated with productivity gains. Taken together, it is clear that unobservables or private information held by the middle managers drive the pattern in the heterogeneity results: that recommended supervisors gain more in terms of retention while gaining little in terms of productivity.

Targeting Allocation Using Skill Scores. Suppose the firm is considering large scale

[28] Specifically, to shut down the effects of unobservables, we set the selection term in second stage to 0.

adoption of the STITCH training, and undertakes a pilot of the training to decide whether the investment is worth it. If the firm pilots by treating the supervisors of half the lines based on nominations of middle managers, as is indeed the usual approach in contexts like this, our results suggest that they would see negligible productivity effects and likely not adopt the program. Given the large average gains we document (and the large returns to investment we document in Section 3.6), this would be a costly error due to decentralization of the allocation decision.

This raises the question of whether there is a way to use the information that middle managers have about the supervisors without fully decentralizing the result. We consider an alternate allocation rule that uses information gleaned from the middle managers about the skill stocks of workers. We focus on two simple and ex-ante reasonable allocation rules: allocate training to production lines with supervisors who have the lowest average baseline score for (1) management skills and (2) industrial engineering skills as elicited by the middle managers.[29] Appendix Figure C.6 presents the results treatment effects. As expected from previous analysis, these allocation schemes substantially outperform random assignment in terms of productivity gains. With half the lines treated, allocation based on management or industrial engineering skills leads to approximately 93% and 88% larger average treatment effects, respectively (ATE of 5.2 pp and 4.9 pp, compared to the ATE of 2.8 pp under randomization). These result suggest that just eliciting information from the middle managers and making allocation decisions based on the elicited information is preferable to full decentralization or randomization in this context. However, this allocation scheme would likely not be incentive compatible over time. Middle managers can learn that elicited information is used for allocating training and distort the information they provide, undoing the informational content of their answers.

3.5.2 Alternative Explanations

Preferred Interpretation: Misaligned Incentives. Our interpretation of the results is that middle managers are targeting retention above and beyond productivity gains due to misaligned incentives between supervisors and the firm. While the key objective of the firm is to target productivity as it is the first order predictor of profits, middle managers face personal costs when supervisors quit and therefore have private incentives to target supervisor retention beyond its implications for productivity. Below, we lay out the key building blocks needed for this interpretation and discuss the evidence.

[29] We consider these allocation rules as ex-ante reasonable because the STITCH training explicitly focuses on the management and industrial engineering skills we emphasize in our survey such as "assigning workers to operations", "meeting targets", and "line balancing." In contrast, technical sewing skills are not related to the content of the STITCH training.

1. Middle managers need to possess relevant information about the supervisors' potential treatment effects to purposefully allocate training based on the incentives they face. Section 3.5.1.1 and Section 3.5.1.3 provide evidence that middle managers possess useful information about baseline stock of skills that can be utilized to allocate training to successfully target productivity gains. Further, Figure 3.1b shows that highly recommended supervisors in the control group have higher future quitting rates, consistent with the idea that middle managers possess information that predicts which managers are at greatest risk of quitting, allowing for the targeting of recommendations towards preventing quitting among supervisors.

2. The supervisors with high middle manager recommendations should respond to treatment more on the retention margin such that aligning recommendations with the risk of quitting is a valuable strategy for middle managers. This is clearly evidenced by the heterogeneous retention impacts of treatment we document in Section 3.4.3.1. Treated supervisors are 28% less likely to quit among high-recommendation supervisors, while there is no discernible effect for low-recommendation supervisors.

3. The pattern of heterogeneity in treatment effects for productivity and retention must be discernible different from each other. That is, if productivity gains and retention gains were strongly positively correlated, we would not be able to discern the relevant margin the middle managers were targeting. However, as discussed in Section 3.4.2.1, high recommendation supervisors (who have high retention gains from training) exhibit negligible productivity gains from training. This negative correlation between retention and productivity gains from training allows us to observe which margin is being targeted by the middle managers.

4. There should be evidence of incentives for the middle managers to target retention above and beyond their implications for productivity. The survey evidence reviewed in Section 3.2.1.3 shows that middle managers do indeed have private incentives to minimize the retention of supervisors due to personal costs from having to substitute for attrited supervisors, search for replacements, and onboard and support new hires.

5. For the results to be consistent with misaligned incentives, the middle manager allocation needs to be suboptimal from the firm's perspective. While the firm would value retention of supervisors as turnover is likely associated with explicit costs and possibly reduced productivity, two pieces of evidence nevertheless show that the observed middle manager allocations are suboptimal for the firm. First, our productivity impacts are ITT at the line level and therefore incorporate any productivity impacts medi-

ated through changes in retention. Therefore, the negligible productivity gains for the highly recommended supervisors are not complemented by any potential productivity gains through the increase in retention. Second, we undertake a calculation in Section 3.6 below that indicates that the cost of an average supervisor's turnover needs to be implausibly large (9 to 18 times the average annual earnings of a supervisor in our study) for the retention gains from the middle manager allocation to be optimal for the firm despite the foregone gains in productivity.

We next discuss additional alternative explanations and why they are unlikely to be the driving factors behind the observed middle manager recommendation patterns.

Discrimination and Favoritism in Middle Manager Recommendations. One explanation for the negative relationship is that middle managers discriminate based on demographics or favoritism. If the characteristics middle managers discriminate on are negatively related to treatment gains, we could observe the negative relationship we see in the data. While subtle forms of discrimination or favoritism would be indeed hard to capture, we do not see strong evidence of discrimination/favoritism in our data. Many demographic characteristics (gender, age, caste, etc.) and measures that may relate to favoritism (coincident tenure, whether the supervisor and the middle manager started the firm in the same year) are included in the LASSO exercise. The only variables related to demographics that are selected in this analysis are whether the supervisor is of general caste (associated with lower recommendation) and whether the supervisor is from out of state (associated with higher recommendation).

In Appendix Table C.11, we directly look for evidence by regressing the middle manager recommendation on many demographic and coincident tenure related characteristics. With all supervisors included, across the 15 covariates, only whether the supervisor's native language is Kannada (native language of the region) is significantly and negatively related to middle manager recommendation. The R^2 of the model is only 1.6% and the joint F-statistic is 1.36 ($p = 0.17$). The second column restricts the sample to the supervisors who are in our productivity analysis sample, a similar restriction to the LASSO exercise. The only variable with a statistically significant (and positive) relationship is whether the supervisor is from out of state, consistent with the LASSO exercise. The R^2 is still low (2.5%) and the joint F-statistic is insignificant ($p = 0.36$). Overall, the available measures in our data do not indicate discrimination or favoritism as clear ulterior motives driving the recommendations.

Rewarding or Hoarding Productive Supervisors. An alternative view is that middle managers might view the training program as a reward for supervisors who are performing well. However, such a selection would presumably lead to a positive correlation between either baseline productivity or the skill scores of the middle managers. None of the three skill scores show up in our LASSO analysis as a good predictor of the middle manager recommendations. Baseline productivity of the supervisor's line does show up in this analysis, but is (weakly) negatively related to the recommendation. These patterns provide no evidence for rewarding productive supervisors through training allocation. In Appendix Table C.12, we present results from regressing the middle manager recommendation on the three skill scores and supervisor's motivation to improve, as elicited from the middle managers. The first two columns show results for the full sample. First, the three skill scores alone explain an exceedingly small fraction of the variation in middle manager recommendation with $R^2 = 0.5\%$. Second, after controlling for supervisor's motivation to improve, which is positively correlated with middle manager recommendations consistent with the LASSO analysis, there is a negative relationship between the middle manager recommendation and the technical skills of the supervisors. This is contrary to what we would expect to see if middle managers aimed to reward good supervisors with the training.[30]

Conversely, if middle managers believe that trained supervisors are more likely to get promoted and stop being a supervisor on the floor, they may recommend low-skill supervisors for training to hoard talent, similar to the mechanism explored in Haegele (2022b). This is a harder interpretation to dismiss given the negative relationship between middle manager recommendation and baseline productivity and baseline technical skill score. However, these relationships are relatively weak. That is, the fact that none of the three skill scores show up as a predictor of rankings in the LASSO makes it unlikely that hoarding motive is the primary driver of the recommendations.

One possible issue is that middle managers may be strategically misreporting the skill scores of the supervisors under them, either to justify their recommendations or to misidentify talented workers to avoid detection of hoarding. Two facts go against this possibility. First, the middle managers are asked about the supervisor skill scores before they are told about the training and asked about its allocation, making it unlikely that the training allocation is inducing them to misreport. Second, as we discuss in Section 3.5.1.1, the skill scores are meaningfully related to both baseline productivity and gains from training, making it unlikely that there is large scale strategic mis-reporting. Finally, in our context, it is unlikely that the middle managers believed the trained supervisor would leave the floor due to a promotion as such promotions are rare for supervisors in this firm.

Mistaken Beliefs About the "Production Function" of Training. In Section 3.5.1.1, we argue that supervisors possess private information about supervisor skills that, if used

[30]For the subsample included in the productivity analysis (columns 3 and 4), this negative relationship between middle manager recommendation and technical score vanishes.

properly, can be utilized to allocate training effectively. Specifically, allocating training to supervisors who have been indicated by their middle managers to have a low level of baseline managerial and industrial engineering skills beats randomization in terms of productivity gains. However, this does not rule out the possibility that middle managers have information about their supervisors, but they systematically misunderstand the production function of the training. For example, supervisors may believe that training is a complement to baseline stock of skills, instead of a substitute. Misallocation would then result not from a lack of knowledge about supervisors themselves, but about the training. While possible, we note that in a world where such a misunderstanding is the main mechanism behind the negative relationship between recommendation and productivity gains, we would expect there to be a strong relationship between supervisor skill scores (specifically for skills that middle managers believe are complementary to training gains) and recommendation. As discussed above, we do not observe this strong relationship.

Multiple Types of Middle Managers. It is possible that there are multiple types of middle managers ranking supervisors in countervailing ways. Some middle managers may be recommending high skill supervisors to reward them while others may be recommending low skill supervisors so as to hoard the high skill supervisors, with a null average relationship. This is a difficult possibility to entirely refute, but we look for evidence that more than one approach or strategy is being employed in the data. We repeat our LASSO analysis of observable determinants of middle manager recommendations a 1000 times on random subsamples of middle managers and their reporting supervisors to see if there exists groups of middle managers that seem to employ different strategies or, in particular, have opposite relationships between key observable determinants such as skill ratings and training recommendations. Overall, we do not find much evidence to support this notion. The details of our approach and results are in Appendix Section C.6.

3.6 Returns on Investment

Finally, we quantify the profit and rate of return to the firm from STITCH training. To do so, we first report treatment effects on supervisor salaries and performance-based incentive payments which inform the indirect costs of training to the firm. We then calculate that STITCH training leads to significant net profit gains of around $4.5 million for the firm. We conclude by providing a back-of-the-envelope calculation that suggests that the cost of supervisor turnover needs to be implausibly large for the middle manager allocation of training to be optimal from the perspective of the firm. This further bolsters our conclusion that decentralized middle manager allocation would be extremely suboptimal for the firm due to agency issues.

Additional Treatment Effects for Cost Calculation. To assess the net profit for the firm, we need to take into account any indirect cost increases associated with the training due to changes in payments to employers. To do so, we assess whether STITCH impacted the salaries of treated supervisors and whether the productivity impacts lead to increased incentive payments which the firm pays out to employees on the basis of performance. First, we find that salary growth of treated supervisors is 0.8 pp higher than the control group, corresponding to 6% of the control mean. Second, treatment increases the likelihood that a production line receives any incentive payment by 3 and 4 pp during and after treatment, respectively, corresponding to 38% and 51% of the baseline means (significant at 10 %). On the intensive margin, we find 26% (not significant at 10%) and 37% (significant at 10%) increases in incentive payments during and after training, respectively. These effects accrue not only to trained supervisors but also to non-supervisory workers on the line. For brevity, details of these analyses are presented in Appendix Sections C.5.9 and C.5.10.[31]

Net Profit and Rate of Return. We quantify the profit and rate of return to the firm from STITCH training by combining our effect estimates on productivity, wage growth, and incentive payments with program cost data and inputs from the accounting department of the firm. Table 3.7 presents our net profit and rate of return calculations.

On the benefit side, we exclusively focus on the 480 production lines included in our main productivity analysis sample. This is conservative as it implicitly assumes the gains for the lines not included in this sample is 0.[32] We combine our productivity effect estimates with the target quantity information for each line-day and revenue/profit margin we obtained from the firm.[33] Overall, counting flow benefits up till six months after training end, the NPV of the benefits from additional productivity is more than $4.5 million. On the cost side, we consider direct costs of the program, such as trainer salaries, equipment, food, and

[31] In addition for completeness, we check for but do not find any effects on supervisor attendance, worker attendance, and worker retention. We caution that lack of worker level results may partly reflect noisy worker-line matches. Details of these analysis can be found in Appendix Sections C.5.11 and C.5.12.

[32] We also do not consider spillovers in this analysis. This is also likely conservative as available evidence is suggestive of some positive spillovers after training.

[33] Consistent with our context, we assume the firm can sell its additional production without changing prices within the neighborhood of current production. In fact, the marketing team routinely overbooks capacity by roughly 20%, because profit is most dramatically impacted when lines lay vacant. The firm has been steadily and consistently expanding capacity for decades with existing factories adding lines and new factories opening, but the firm is careful only to expand capacity to accommodate predictable and steady excess demand. The loss is asymmetric, given that labor makes up 30% of production costs but profit margins are only 5%. That is, the firm would rather shift 20% of orders into the future (or even turn down orders as they routinely do) than risk having unutilized capacity.

program development costs ($13,085), costs associated with additional incentive payments for lines with treated supervisors ($31,976), and increased salary of treated supervisors ($ 36,815).[34] Overall, we estimate the total costs to be around $82,000. The net profit from the program considering cost and benefit flows up through 6 months after training end is $4,460,669 (corresponding to $4,807 per overall treated supervisor and $10,302 per treated supervisor working at one of the 480 analysis lines). The net rate of return is thus around 54 times the training cost.[35]

How Much Retention Needs to Cost for the Firm to Prefer Middle Manager Allocation. Finally, we undertake a simple back-of-the-envelope calculation to give a rough estimate of how much additional supervisor turnover would need to cost (above and beyond its productivity effects which are already internalized in our line-level intent-to-treat productivity impact estimate) for the middle manager allocation to be justified from the perspective of the firm. To do this, we compare the avoided supervisor turnover to the missed productivity gains that would result from the middle manager allocation.

On the retention side, we note that around 35% of the supervisors in the control group leave the firm by the end of our period.[36] Applying this percentage, we assume 325 of the 928 supervisors in the treatment group would leave the firm absent treatment. The quitting hazard ratio for treated among the highly recommended supervisors, presented in Appendix Table C.14, is 28% (while it is almost 0 for low treatment supervisors). Using this point estimate, we conclude that the middle manager allocation would avoid 90 of the 325 turnovers that would take place absent training. We conservatively assume that the middle manager allocation would lead to half the productivity benefits that we observe from the random allocation, while keeping the costs the same.[37] This implies that each additional turnover would need to cost the firm approximately $24,800 for the middle manager allocation to be profitable from the firms perspective, as compared to the random allocation.[38] In the extreme

[34]Because the training took place on Sundays, the off day of supervisors, costs associated with lost production hours are not a part of our cost calculation.

[35]Note however the cost calculations do not include the cost of the authors' expertise in conducting the prior study (Adhvaryu et al., 2022d) which informed the curriculum. One could argue that prior to this study the cost of acquiring the content for this curriculum would be extremely high for the firm.

[36]This value only includes supervisors who were present at the start of the training.

[37]In fact, we know that the productivity gains are heavily concentrated in lines with lower average middle manager recommendations, so the productivity benefits would likely fall substantially more.

[38]The firm reports that quantifying the cost of replacing workers who leave is in general difficult. Though they had no such calculation for supervisors, their best assessment of the cost of replacing a machine operator was roughly 20,000 INR which amounts to roughly 300 USD at the time of the study. This is less than 1.5% of the cost needed to justify the middle manager training allocation with respect to returns to the firm. The cost of replacing supervisors would, if anything, be less given that machine operators are mostly recruited from distant villages, trained for several months and relocated at the firm's expense to the city, while supervisors are generally hired from other nearby factories or promoted from within the firm.

Table 3.7: Return on Investment Calculations for 6 Months After Program End

Total Benefit (Only For the 480 Sewing Lines Included in Analysis Sample)	$4,542,544
Additional Productivity (Lines with STITCH Trained Supervisors)	$4,542,544
Total Cost	-$81,875
STITCH Training Cost (Development, Trainer Salary, Materials, and Refreshments)	-$13,085
Additional Incentive Payments (Lines with STITCH Trained Supervisors)	-$31,976
Increased Salary (STITCH Trainees)	-$36,815
Net Benefit	**$4,460,669**
Net Rate of Return	**54X**
Assumptions	
Revenue per Additional Garment	$7
Profit Margin on Revenue from Additional Productivity	20 %
Interest Rate	10 %
Exchange Rate (INR per 1 USD)	65

Note: All values in April 2017 present values. Productivity calculations only covers the 480 lines included in our main analysis sample. Period of interest is from training start to 6 months after training end for each factory, consistent with our analysis. Additional garments due to training is calculated by multiplying the average target quantity for a given line-month with the relevant coefficient (based on during/post training) and the fraction of line supervisors treated. We then assume there are 25 production days on a given month. Revenue per additional garment is taken from the accounting department of the firm. Profit margin on revenue from additional productivity is calculated as 80% of the percent labor contribution to cost (25%) as guided by the accounting office. For each line-day, we find the treatment effect coefficients on incentive payments are 26 INR and 16 INR for during/after training. We add this cost multiplied by fraction of supervisors treated for each line-day for the included lines to get additional incentive payments. For increase in salary, we multiply treatment effect on salary growth (0.008) with average baseline salary of supervisors in April 2017 (14,855 INR) with the number of trained supervisors for each month. Observe that this is conservative as it includes all trained supervisors (not just the ones for the analysis lines) and assumes percent increase in salaries take place immediately. Materials and food cost amounted to 30,000 INR and 150,000 INR respectively. Cost of development was 70,500 INR. Trainer salaries were 600,000 INR. Exchange rate is the rate in the beginning of April 2017.

case of approximately zero productivity gains under middle manager allocation (which is not unrealistic given the patterns we document above), each supervisor replacement would need to cost the firm around $49,600 to make the middle manager allocation better than random allocation from the firm's perspective. To put this in perspective, these costs are roughly 9 to 18 times the average annual earnings of a supervisor in our study, which are around $2,700 at the start of the training. Given the large magnitudes, it is implausible that the additional retention from the middle manager allocation would make up for the foregone productivity gains from the firm's perspective. Note that these calculations compare the middle manager recommendation allocation to random assignment, but the middle manager recommendation assignment rule is even more starkly suboptimal when compared to an alternative assignment rule which outperforms random assignment (e.g., one that uses baseline skill deficiencies elicited from the middle managers as discussed in section 3.5.1.3).

3.7 Conclusion

A recent empirical literature has documented the value of having multiple layers of management (Caliendo and Rossi-Hansberg, 2012; Caliendo et al., 2015, 2020) and decentralizing responsibilities and decisions to lower levels of the hierarchy (Bloom and Van Reenen, 2011; Bloom et al., 2014; Aghion et al., 2021). These studies argue that middle managers may have some private information and/or specialized understanding that makes them better equipped for making decisions; however, the classic tradeoff is that this decentralization creates a principal-agent structure in which the middle manager may act according to private incentives which do not align perfectly with those of the organization and that limited information at the top of the organization may make enforcing organizational incentives difficult (Acemoglu et al., 2007; Aghion et al., 2014).

To study this exact tradeoff as it relates to the allocation of managerial training within a firm, we elicited from middle managers rankings of which line supervisors should be prioritized for training and then randomized access to training within these rankings. We find that line supervisors gained substantial knowledge from the training and productivity of teams managed by trained supervisors increased substantially and persistently on average. However, these productivity gains were quite heterogeneous, with line supervisors recommended highly by middle managers to receive the training actually gaining little to nothing from the training in terms of productivity.

On the other hand, training generated a significant positive treatment effect on retention, with these impacts driven entirely by the highly recommended supervisors. In addition high recommendation supervisors in the control group were more likely to quit in the absence of training than were low recommendation control supervisors. We adapt a recent approach by Dal Bó et al. (2021) to decompose the allocation decisions of middle managers into observable and unobservable components. This analysis confirms that substantial variation (at least 80%) in middle manager recommendations derives from unobserved drivers, and that this unobserved component (perhaps most indicative of the private information to be leveraged via decentralization of the training allocation decision) positively predicts improvements in retention despite negatively predicting productivity gains.

Taken all together, the results suggest that middle managers may know which supervisors are most likely to quit and that allocating a training investment of this sort to them may improve their retention. Accordingly, middle managers appear to tailor their recommendations to take advantage of this potential improvement in retention. We note that the return on investment implied by these net productivity gains is several orders of magnitude larger than any monetary costs borne by the firm to screen and train new supervisors. Accordingly, the firm clearly favors allocating the training to maximize gains in productivity (as would workers and supervisors who all earn significantly greater incentive pay as a result of treatment effects on productivity), but the middle managers have competing incentives to improve line supervisor retention in order to minimize the private burden to them of screening and training replacements and covering the supervisor duties in the interim.

Importantly, the retention of line supervisors which middle managers appear to prioritize is, of course, not without value or importance to the firm, but rather the firm would simply prioritize productivity gains (which deliver orders of magnitude larger returns) when the two priorities are at odds, as turns out to be the case in our scenario. Our results show that though the average productivity gains from a random allocation were large, persistent, and generated tremendous return on investment, if the supervisors who gained little to nothing had been targeted (as would have been the case if the allocation decision were decentralized to middle managers) the gains and return on investment would have been negligible.

Indeed, we note that the very design of the trial reported on in this paper was motivated by anecdotal conversations with upper management at the firm regarding how investments like the one we evaluate get piloted and rolled out in the firm. These conversations revealed that many such programs are proposed and considered over the course of the year, often from buyers with whom the firm wants to maintain a strong relationship (Adhvaryu et al., 2020). Given these programs are costly particularly in terms of time and effort for their coordination and implementation, the firm often pilots these programs with a subset of production lines or workers before deciding to roll them out across the entire firm.

The most likely way these pilot lines and workers are selected is via a decentralized recommendation much like the one we elicited in the study. Accordingly, if the firm were to undertake exactly this pilot approach with respect to the program we study here, we note that they would have believed the gains to be null and would have aborted the program after the pilot, forfeiting $4.5M in gains. In this sense, our results provide one potential explanation for why managerial quality remains low on average and highly varied in many firms despite growing academic evidence of potentially large gains from investments in management such as the training program we evaluate. That is, when returns to such investments are heterogeneous (due to heterogeneity in baseline stocks of these skills and productivity across supervisors and teams within the firm) and allocation of costly resources is decentralized (potentially as a means of piloting to inform investment decisions), investment decisions may be made on the basis of inaccurate estimates of returns leading to underinvestment.

APPENDIX A

Appendix to Chapter I

A.1 Appendix Figures and Tables

Table A.1: Summary Statistics

	Mean	SD	Percentiles 1st	25th	50th	75th	99th
Migration Rate	0.34	0.26	0.00	0.14	0.28	0.48	1.11
Average Annual Migrant Wages (USD)	5371.17	1225.88	3736.68	4598.98	5140.51	5884.59	9182.36
Average Contract Duration (Months)	22.83	1.59	17.52	22.05	23.08	23.91	25.78
Share Migrating for Occupation Quartile:							
1st (Lowest Paying)	0.67	0.17	0.23	0.56	0.69	0.79	1.00
2nd	0.15	0.11	0.00	0.08	0.13	0.21	0.50
3rd	0.07	0.06	0.00	0.03	0.06	0.10	0.25
4th (Highest Paying)	0.10	0.08	0.00	0.05	0.09	0.14	0.36
Share Migrating to Country Quartile:							
1st (Lowest Paying)	0.77	0.14	0.35	0.69	0.78	0.86	1.00
2nd	0.02	0.02	0.00	0.00	0.01	0.02	0.11
3rd	0.09	0.08	0.00	0.04	0.08	0.13	0.33
4th (Highest Paying)	0.12	0.10	0.00	0.06	0.10	0.16	0.47
Migrant Cohort Demographics:							
Average Age	31.81	1.69	27.11	30.90	31.87	32.79	36.00
Share Male	0.36	0.18	0.00	0.23	0.34	0.48	0.83
Migrant Stock Education:							
Share Completed High School	0.85	0.14	0.30	0.83	0.90	0.94	1.00
Share Completed College	0.39	0.14	0.04	0.29	0.38	0.49	0.71
Working Age Population Education:							
Share Completed High School	0.57	0.17	0.18	0.45	0.57	0.69	0.97
Share Completed College	0.14	0.06	0.03	0.10	0.13	0.17	0.33

Notes: Summary statistics for municipality-year covering 2007 to 2016. Education estimates for the migrant stock and working age population (aged 20 to 64) from 2007, 2010, and 2015 population censuses. The remaining estimates are from the administrative contract data.

Table A.2: Top 20 Migration Destination Countries (2007-2016)

Destination	Number of Migrants	Percent of Migrants	Percent of Migrants in Destination that are Filipino
Saudi Arabia	1404274	35.3	4.8
United Arab Emirates	542319	13.6	6.5
Qatar	378412	9.5	8.8
Kuwait	337952	8.5	6.3
Taiwan	332949	8.4	NA
Hong Kong	264421	6.7	4.2
Singapore	113154	2.8	.6
Malaysia	74900	1.9	3.4
Bahrain	70490	1.8	6.3
Japan	57167	1.4	10.3
Oman	51640	1.3	1.9
Canada	49552	1.2	7.9
South Korea	38703	1	3.8
Brunei Darussalam	28927	.7	13.1
Papua New Guinea	19169	.5	4.7
Jordan	18474	.5	.1
Italy	18155	.5	2.5
Israel	15409	.4	NA
United States	15194	.4	4.1
Cyprus	14543	.4	3.9
Other	130423	3.3	

Notes: Author's calculations from the administrative migration data for columns 1 and 2. Migrant percentage is based on 2015 UN Department of Economic and Social Affairs estimates for column 3.

Table A.3: Alternative Measurement

	(1)	(2)	(3)	(4)
	Migration Rate...			
	per 2000 pop	per 2007 pop	ln(migrants)	ln(mean wage)
$T_{m,[t,t-1]}$ ($\beta_{ShortRun}$)	1.405***	1.412***	0.044**	-0.014***
	(0.503)	(0.449)	(0.019)	(0.005)
	[0.433]	[0.396]		[0.005]
$T_{m,[t-2,t-3]}$ ($\beta_{MediumRun}$)	1.475***	1.459***	0.029*	-0.006
	(0.478)	(0.423)	(0.015)	(0.004)
	[0.438]	[0.392]		[0.004]
Observations	15,970	15,970	15,970	15,788
Adjusted R2	0.894	0.892	.	0.874
Mean Dep. Var.	39.260	34.581	231.675	5.573

Notes: Unit of observation is municipality-year. All regressions include municipality fixed effects. Migration rates are in per 10,000 population. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. *** p<0.01, ** p<0.05, * p<0.10

Table A.4: Typhoons Increase Migration and Decrease New Migrant Wages - Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
	Migrants per 10,000 capita			Mean ln(wage)			25th Percentile ln(wage)			50th Percentile ln(wage)			75th Percentile ln(wage)		
$T_{m,[t,t-1]}$ ($\beta_{ShortRun}$)	1.078***	1.244***	1.111***	-0.008***	-0.010***	-0.011***	-0.006**	-0.010***	-0.008***	-0.013***	-0.016***	-0.012***	-0.016***	-0.013**	-0.018***
	(0.332)	(0.413)	(0.397)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)	(0.003)	(0.004)	(0.005)	(0.005)	(0.005)
	[0.359]	[0.413]	[0.356]	[0.002]	[0.002]	[0.002]	[0.003]	[0.003]	[0.002]	[0.004]	[0.004]	[0.003]	[0.004]	[0.004]	[0.004]
$T_{m,[t-2,t-3]}$ ($\beta_{MediumRun}$)	1.391**	1.463***	1.338***	-0.001	-0.005***	-0.007***	-0.004	-0.006**	-0.005**	-0.013***	-0.010***	-0.009**	-0.003	-0.005	-0.009**
	(0.333)	(0.492)	(0.400)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)	(0.002)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)
	[0.525]	[0.433]	[0.400]	[0.002]	[0.002]	[0.002]	[0.002]	[0.003]	[0.002]	[0.003]	[0.003]	[0.003]	[0.004]	[0.003]	[0.003]
Year FE	Yes	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes	No	No
Year-IslandGroup FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Base. Char. Trend	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
Observations	15,970	15,970	15,970	15,788	15,788	15,788	15,768	15,768	15,768	15,768	15,768	15,768	15,768	15,768	15,768
Adjusted R2	0.969	0.877	0.886	0.902	0.904	0.906	0.659	0.676	0.691	0.830	0.837	0.848	0.850	0.863	0.866
Mean Dep. Var.	230.659	33.644	33.644	5.476	5.476	5.476	5.287	5.287	5.287	5.377	5.377	5.377	5.576	5.576	5.576

Notes: Unit of observation is municipality-year. All regressions include municipality fixed effects. Migration rate is calculated per 10,000 capita. Observation numbers for columns 4-15 are lower due to municipality-years with no migration. In columns 4-15, observations are weighted by the number of migrants making up each cell. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. *** p<0.01, ** p<0.05, * p<0.10

Table A.5: Typhoons Increase Migration and Decrease New Migrant Wages - Winsorized Typhoon Exposure

	(1)	(2)	(3)	(4)	(5)
		Migrant Wages ...			
	Migration rate	mean ln(wage)	ln 25th pct.	ln 50th pct.	ln 75th pct.
$T_{m,[t,t-1]}$ ($\beta_{ShortRun}$)	1.470***	-0.011***	-0.011***	-0.017***	-0.014**
	(0.405)	(0.003)	(0.003)	(0.003)	(0.005)
	[0.425]	[0.003]	[0.003]	[0.004]	[0.005]
$T_{p,[t-2,t-3]}$ ($\beta_{MediumRun}$)	1.678***	-0.005**	-0.007**	-0.011***	-0.005
	(0.525)	(0.002)	(0.003)	(0.004)	(0.004)
	[0.464]	[0.002]	[0.003]	[0.003]	[0.004]
Observations	15,970	15,788	15,768	15,768	15,768
Adjusted R2	0.877	0.904	0.677	0.837	0.863
Mean Dep. Var.	33.644	5.476	5.287	5.377	5.576

Notes: Unit of observation is municipality-year. All regressions include unit and year-by-island-group fixed effects. Typhoon exposure is winsorized at 99%. Migration rate is calculated per 10,000 capita. Observation numbers for columns 2-5 are lower due to municipality-years with no migration. In columns 2-5, observations are weighted by the number of migrants making up each cell. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered SEs.

Table A.6: Occupation and Destination Country Quartile Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Share of Migrants Going To:							
	Country Quartiles				Occupation Quartiles			
	1st	2nd	3rd	4th	1st	2nd	3rd	4th
$T_{m,[t,t-1]}$ ($\beta_{ShortRun}$)	0.014***	-0.001	-0.004*	-0.009***	0.008**	-0.007***	-0.001	0.000
	(0.003)	(0.001)	(0.002)	(0.003)	(0.003)	(0.002)	(0.001)	(0.001)
	[0.003]	[0.000]	[0.002]	[0.002]	[0.003]	[0.002]	[0.001]	[0.001]
$T_{m,[t-2,t-3]}$ ($\beta_{MediumRun}$)	0.018***	-0.000	-0.004**	-0.013***	-0.004*	-0.001	0.002***	0.002**
	(0.003)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
	[0.003]	[0.000]	[0.002]	[0.002]	[0.002]	[0.002]	[0.001]	[0.001]
Observations	15,785	15,785	15,785	15,785	15,785	15,785	15,785	15,785
Adjusted R2	0.825	0.521	0.772	0.770	0.931	0.881	0.776	0.825
Mean Dep. Var.	0.768	0.016	0.091	0.121	0.673	0.153	0.070	0.102
SD Dep. Var.	0.141	0.025	0.075	0.101	0.166	0.107	0.060	0.077

Notes: Unit of observation is municipality-year. All regressions include municipality and year-by-island-group fixed effects. Observations are weighted by number of migrants making up each cell. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered standard erros.

Table A.7: Additional Migrant Level Results

	(1)	(2)	(3)	(4)	(5)	(6)
	ln(wage)			$\mathbf{1}$[Construction]	$\mathbf{1}$[2nd occ. quartile]	
$T_{m,[t,t-1]}$ ($\beta_{ShortRun}$)	-0.010***	-0.003***	-0.000	-0.002	-0.006***	-0.005***
	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
$T_{p,[t-2,t-3]}$ ($\beta_{MediumRun}$)	-0.005***	-0.002**	0.000	0.001	-0.000	-0.001
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$\mathbf{1}$[Construction Occ]						0.750***
						(0.003)
Ctry by Occ FE	No	Yes	Yes	No	No	No
Ctry by Occ by $\mathbf{1}[GR_t]$	No	No	Yes	No	No	No
Observations	3,637,967	3,637,240	3,636,816	3,663,430	3,651,563	3,651,563
Adjusted R2	0.063	0.696	0.704	0.073	0.059	0.450
Mean Dep. Var.	5.527	5.527	5.527	0.119	0.169	0.169
SD Dep. Var.	0.471	0.471	0.470	0.324	0.375	0.375

Notes: Unit of observation is individual migrant contracts. Typhoon exposure index is at the municipality level. All regressions include municipality and year-by-island-group fixed effects. **1**[Construction] is an indicator variable for whether the contract occupation is related to construction. **1**[2nd occ. quartile] is an indicator for whether the contract occupation is in the second wage quartile. $1[GR_t]$ is an indicator for great recession (2008 or 2009). Province clustered standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered SEs.

Table A.8: Typhoons Don't Change the Characteristics of Sampled Households in FIES

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		HH Head Characteristics					
	Household Size	Male	Age	Completed Prim. School	Completed High School	Completed Some College	Completed College
Panel A: Unweighted							
$T_{p,[t,t-1]}$ ($\beta_{ShortRun}$)	-0.011	0.003	-0.011	0.004	0.002	0.001	0.001
	(0.014)	(0.002)	(0.086)	(0.005)	(0.005)	(0.005)	(0.003)
$T_{p,[t-2,t-3]}$ ($\beta_{MediumRun}$)	-0.009	-0.000	-0.138	-0.002	-0.005	-0.001	-0.000
	(0.015)	(0.003)	(0.143)	(0.004)	(0.005)	(0.006)	(0.005)
Observations	306,315	306,315	306,315	306,315	306,315	306,315	306,315
Mean Dep. Var.	4.616	0.786	50.226	0.767	0.468	0.228	0.110
SD Dep. Var.	2.155	0.410	14.376	0.423	0.499	0.419	0.313
Panel B: Weighted by Provided Survey Weights							
$T_{p,[t,t-1]}$ ($\beta_{ShortRun}$)	-0.016	0.003	-0.073	-0.002	-0.002	-0.005	-0.002
	(0.012)	(0.002)	(0.084)	(0.003)	(0.004)	(0.003)	(0.002)
$T_{p,[t-2,t-3]}$ ($\beta_{MediumRun}$)	-0.005	0.001	-0.034	0.000	-0.003	-0.001	-0.001
	(0.015)	(0.002)	(0.149)	(0.002)	(0.003)	(0.003)	(0.003)
Observations	306,315	306,315	306,315	306,315	306,315	306,315	306,315
Mean Dep. Var.	4.616	0.786	50.226	0.767	0.468	0.228	0.110
SD Dep. Var.	2.155	0.410	14.376	0.423	0.499	0.419	0.313

Notes: Household level regression using 2006, 2009, 2012, 2015, and 2018 FIES data. Unit of observation is a household. Typhoon exposure is measured at the province-year level. All regressions include province and year-by-island-group fixed effects. Province clustered standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.10

Table A.9: Typhoons Increase Province Level Migration and Remittance per Capita

	Migration Rate	Mean ln(wage)	ln(Remittance per Capita)
$T_{p,[t,t-1]}$ ($\beta_{ShortRun}$)	2.699***	-0.011***	0.072***
	(0.712)	(0.003)	(0.027)
	[0.657]	[0.003]	[0.027]
$T_{p,[t-2,t-3]}$ ($\beta_{MediumRun}$)	2.193***	-0.005***	0.060*
	(0.699)	(0.002)	(0.034)
	[0.985]	[0.003]	[0.032]
Weights	Population	Cell Size	Population
Observations	790	790	395
Adjusted R2	0.923	0.966	0.915
Mean Dep. Var.	40.318	5.529	7.934

Notes: Unit of observation is province-year. Columns 1-2 uses administrative contract data. Column 3 uses 2006, 2009, 2012, 2015, and 2018 FIES data. Migration rate is per 10,000 capita. All regressions include province and year-by-island-group fixed effects. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets *** p<0.01, ** p<0.05, * p<0.10 based on province clustered standard errors.

Table A.10: Domestic Income Isn't Increasing in Migrant Demand Index

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Total Domestic Income Per Cap.		Wage Income Per Cap.		Entrepreneurial Income Per Cap.		Other Income Per Cap.	
$D_{p,t-1}$	428.0	-237.8	55.5	-190.5	175.5	91.0	122.3	-73.0
	(551.8)	(635.3)	(296.8)	(391.8)	(146.5)	(206.8)	(216.5)	(210.4)
HH Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Counterfactual D	No	Yes	No	Yes	No	Yes	No	Yes
Observations	306,315	306,315	306,315	306,315	306,315	306,315	306,315	306,315
Clusters	79	79	79	79	79	79	79	79
Mean Dep. Var.	44572.0	44572.0	23071.1	23071.1	10435.9	10435.9	12670.7	12670.7
SD Dep. Var.	46730.3	46730.3	32375.8	32375.8	18161.3	18161.3	18428.6	18428.6

Notes: Household level regression using 2006, 2009, 2012, 2015, and 2018 FIES data. Unit of observation is a household. All regressions include province and year-by-island-group fixed effects. Demand index is demeaned within provinces. Observations are weighted by the provided sampling weights. Province clustered standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered standard errors.

Table A.11: Top 20 Export, Import, and FDI Partners for 2007 - 2016

Country	Total export	Percent export	Mig. Share	Country	Total import	Percent import	Mig. Share	Country	Total FDI	Percent FDI	Mig. Share
Japan	94	18.58	1.44	China	74	11.95	.14	Japan	439	24.5	1.44
USA	78	15.38	.38	USA	69	11.12	.38	Netherlands	361	20.18	.01
China	58	11.41	.14	Japan	67	10.88	1.44	USA	272	15.19	.38
Hong Kong	48	9.51	6.65	Singapore	50	8.08	2.85	BVI	125	6.99	0
Singapore	39	7.78	2.85	Taiwan	44	7.07	.1	Korea, Rep.	116	6.47	.97
Netherlands	24	4.78	.01	Korea, Rep.	42	6.77	.97	Singapore	90	5.03	2.85
Korea, Rep.	23	4.58	.97	Thailand	36	5.85	.04	Cayman	57	3.18	.04
Germany	23	4.49	.01	Saudi Arabia	28	4.52	35.32	China	46	2.59	.14
Taiwan	19	3.73	8.37	Malaysia	26	4.23	1.88	Australia	44	2.48	.36
Thailand	18	3.6	.04	Indonesia	26	4.14	.07	UK	30	1.7	.17
Malaysia	14	2.77	1.88	Hong Kong	18	2.85	6.65	Switzerland	23	1.28	.01
Indonesia	6	1.19	.07	Germany	17	2.69	.01	Germany	23	1.28	.01
Vietnam	5	1.04	.06	Vietnam	13	2.08	.06	Taiwan	21	1.15	8.37
UK	5	.9	.17	UAE	11	1.85	13.64	Hong Kong	11	.63	6.65
Australia	5	.9	.36	France	10	1.62	0	Thailand	10	.56	.04
Canada	4	.8	1.25	Australia	9	1.52	.36	India	9	.49	.03
Belgium	4	.79	.01	India	8	1.27	.03	Malaysia	8	.46	1.88
France	4	.72	0	Russia	6	.98	.08	Canada	8	.43	1.25
India	3	.58	.03	Qatar	5	.76	9.52	France	5	.26	0
Mexico	3	.56	.01	New Zealand	4	.68	.22	Denmark	3	.14	0
Others	30	5.92		Others	56	9.08		Others	90	5.03	

Notes: Export and Import data are from COMTRADE. Values in billions $. FDI data from Philipines Statistical Agency reports. Values in billions of real 2010 Phps. Mig share corresponds to the fraction of all migrants in my data (from 2007-2016) going to the relevant destination (in percent)

Table A.12: P-values from Alternative Data Generating Process Assumptions

	(1)	(2)	(3)	(4)	(5)
	Baseline	+ AR(2)	+ country FE	+ gulf	+ AR(2) + ctry FE + gulf
Outcome: Migration per 10,000 capita					
Coefficient on $T_{m,[t,t-1]} \times D_{m,t}$	**1.256**	**1.493**	**2.070**	**1.705**	**1.746**
RI p-val (empirical dist)	0.066	0.028	0.020	0.032	0.025
RI p-val (normal dist)	0.010	0.003	0.000	0.000	0.005
Outcome: ln(mean wage)					
Coefficient on $T_{m,[t,t-1]} \times D_{m,t}$	**0.002**	**0.002**	**0.003**	**0.005**	**0.002**
RI p-val (empirical dist)	0.716	0.287	0.772	0.651	0.845
RI p-val (normal dist)	0.590	0.202	0.637	0.358	0.726

Notes: Each cell corresponds to a different regression and presents the randomization inference p-value for the coefficient on the interaction term of interest (β_1) from estimating equation 1.12. Each column corresponds to an alternative assumption about the DGP underlying the randomization inference procedure, described in A.5.6. Column 2 introduces AR2 term, column 3 introduces country fixed effects, column 4 introduces interaction between year and gulf country dummy, column 5 introduces all three. "Empirical dist." means error terms of re-sampled from the empirical distribution of the error terms within a given year. "Normal dist." means error terms are sampled from a normal distribution matching the mean and variance of the empirical error terms in a given year. Regressions include the mean of the counterfactual demand indices, corresponding to column (2) of Table 1.4.

Table A.13: Occupation and Destination Country Quartile Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Share of Migrants Going To:							
	Occupation Quartiles				Country Quartiles			
	1st	2nd	3rd	4th	1st	2nd	3rd	4th
$T_{m,[t,t-1]}$	0.008***	-0.006***	-0.002**	-0.000	0.010***	-0.000	-0.003*	-0.007**
	(0.002)	(0.002)	(0.001)	(0.001)	(0.003)	(0.000)	(0.002)	(0.003)
	[0.002]	[0.002]	[0.001]	[0.001]	[0.004]	[0.001]	[0.002]	[0.003]
$D_{m,t}$	-0.008	-0.005	0.001	0.012**	0.029***	0.002	-0.020***	-0.009
	(0.008)	(0.007)	(0.003)	(0.006)	(0.007)	(0.002)	(0.005)	(0.006)
	[0.008]	[0.006]	[0.002]	[0.005]	[0.008]	[0.002]	[0.004]	[0.007]
	{0.456}	{0.573}	{0.636}	{0.283}	{0.186}	{0.035}	{0.058}	{0.420}
$T_{m,[t,t-1]} \times D_{p,t}$	-0.016**	0.009***	0.003	0.004	0.019***	0.002	-0.011**	-0.010**
	(0.007)	(0.003)	(0.002)	(0.004)	(0.006)	(0.002)	(0.004)	(0.004)
	[0.006]	[0.003]	[0.002]	[0.003]	[0.008]	[0.002]	[0.005]	[0.005]
	{0.040}	{0.089}	{0.144}	{0.148}	{0.034}	{0.073}	{0.077}	{0.077}
Counterfactual D	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,516	13,516	13,516	13,516	13,516	13,516	13,516	13,516
Clusters	79	79	79	79	79	79	79	79
Mean Dep. Var.	0.667	0.155	0.072	0.104	0.757	0.016	0.095	0.127

Notes: Unit of observation is municipality-year. All regressions include municipality and year-by-island-group fixed effects. Observations are weighted by number of migrants making up each cell. Province clustered standard errors in parenthesis. Standard errors robust to spatial (200 km) and serial (10-year) correlation in square brackets. *** p<0.01, ** p<0.05, * p<0.10 based on province clustered standard erros.

Figure A.1: Temporary International Labor Migrant Flows Across Countries (2006-2016)

(a) Average Total Migrants and Share of 2010 Working Age Population

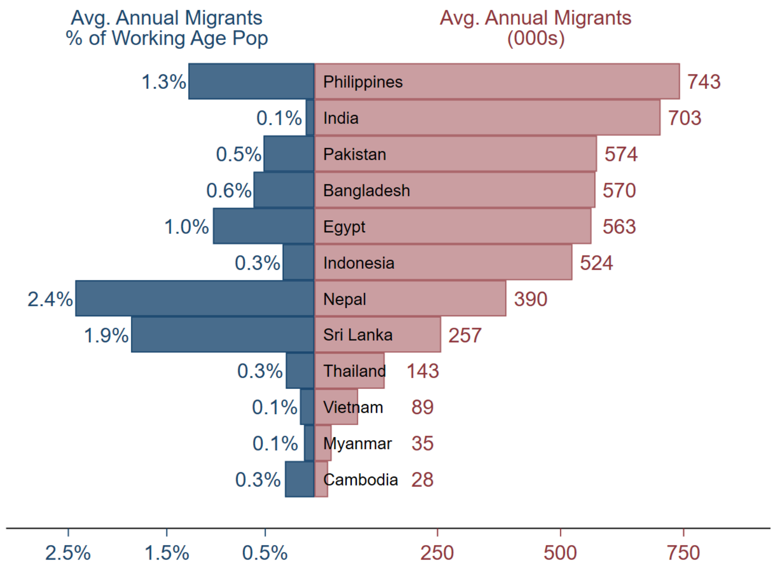

(b) Evolution of Migrant Flows Myanmar, Cambodia, and Vietnam

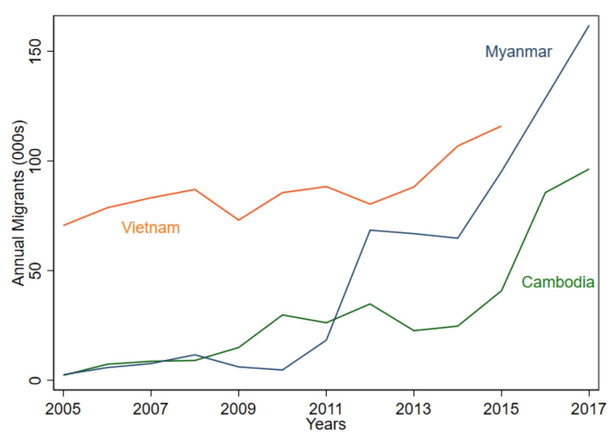

Notes: Data from Philippine Overseas Employment Administration (Philippines); Ministry of External Affairs (India); Bureau of Emigration and Overseas Employment (Pakistan); Bureau of Manpower Employment and Training (Bangladesh); National Board on the Placement and Protection of Indonesian Overseas (Indonesia, through ILOSTAT); Department of Foreign Employment (Nepal); Bureau of Foreign Employment (Sri Lanka); Office of Overseas Employment Administration, Department of Employment (Thailand, through ILOSTAT); Department of Overseas Labour (Vietnam, through ILOSTAT); Department of Labour, Ministry of Labour, Immigration and Population (Myanmar, through ILOSTAT); Economic Census and Department of Employment and Manpower, Ministry of Labour and Vocational Training (Cambodia, through ILOSTAT). Working Age Population is based on 2010 values from the World Development Indicators. Data for 2016 is missing for Thailand, Vietnam, and Myanmar. Data for 2006 and 2007 are missing for Nepal.

Figure A.2: Annual Migration Flows from the Philippines

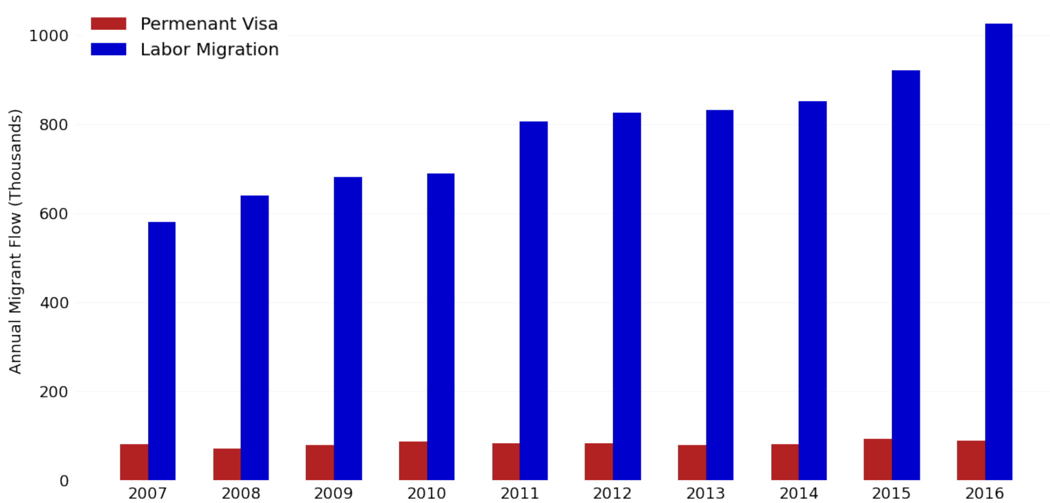

Notes: Data from POEA and CFO reports.

Figure A.3: Municipality Share of Migrants Going to Selected Destinations (2007-2016)

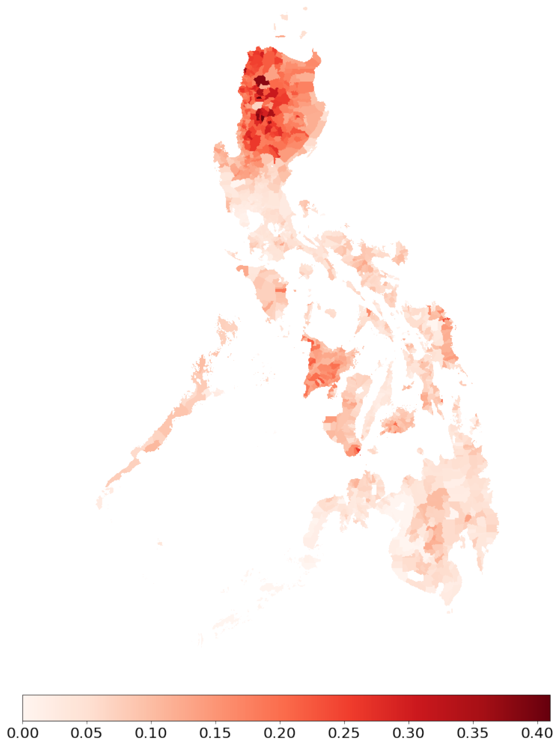

(a) Hong Kong

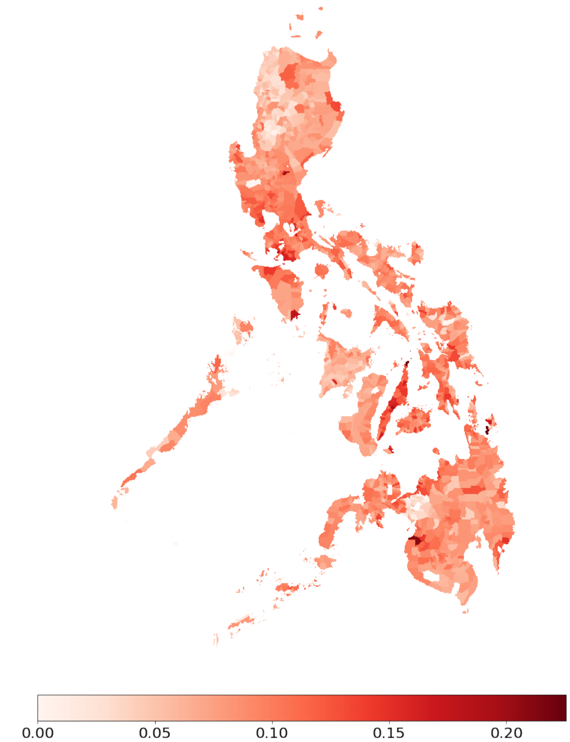

(b) Qatar

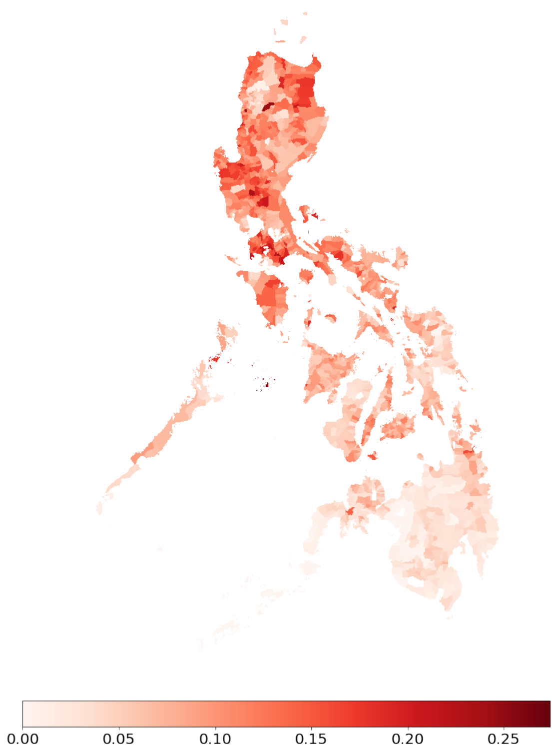

(c) Taiwan

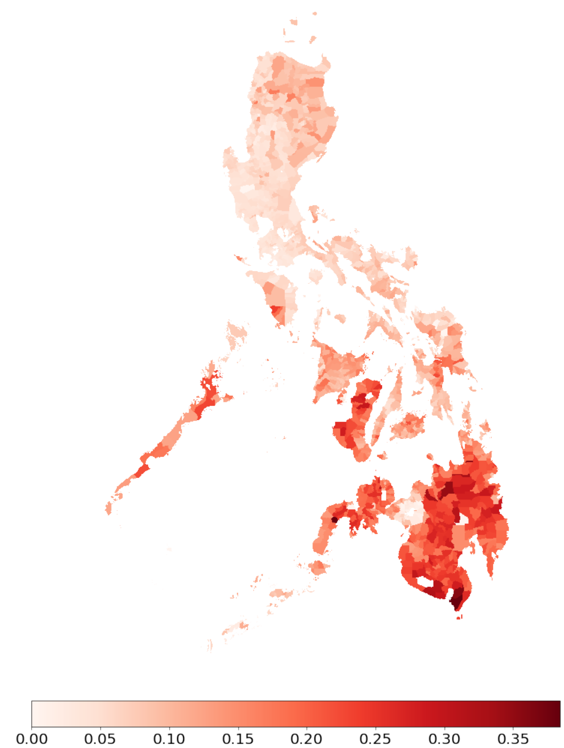

(d) Kuwait

Figure A.4: Robustness: Dropping Provinces One-by-One

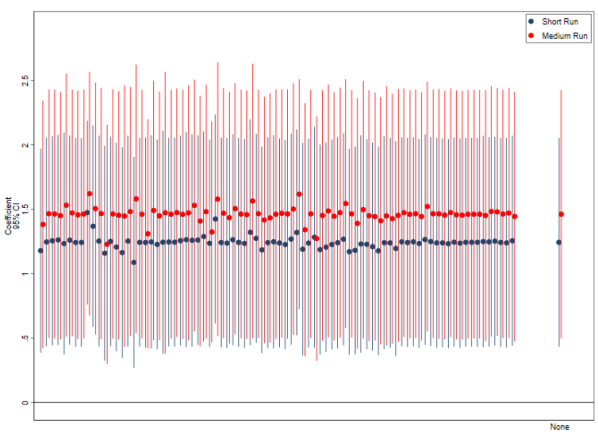

(a) Migration rate

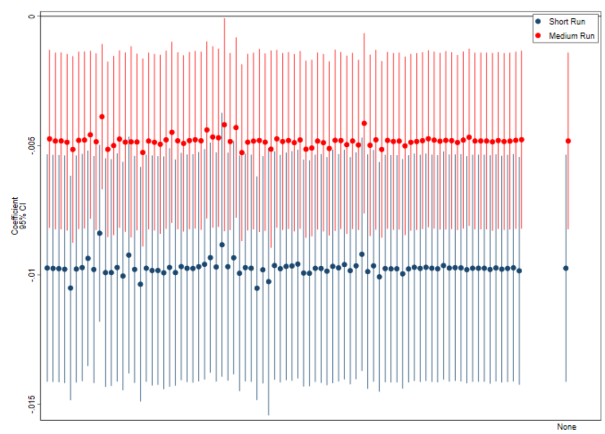

(b) Mean ln(Wage)

Notes: β_{SR} and β_{MR} from summary specification 1.8 are plotted. Provinces are dropped one-by-one from the sample. None refers to estimates with no province dropped. Confidence intervals are based on province clustered standard errors.

Figure A.5: Educational Attainment of Migrant Stock Following Typhoons

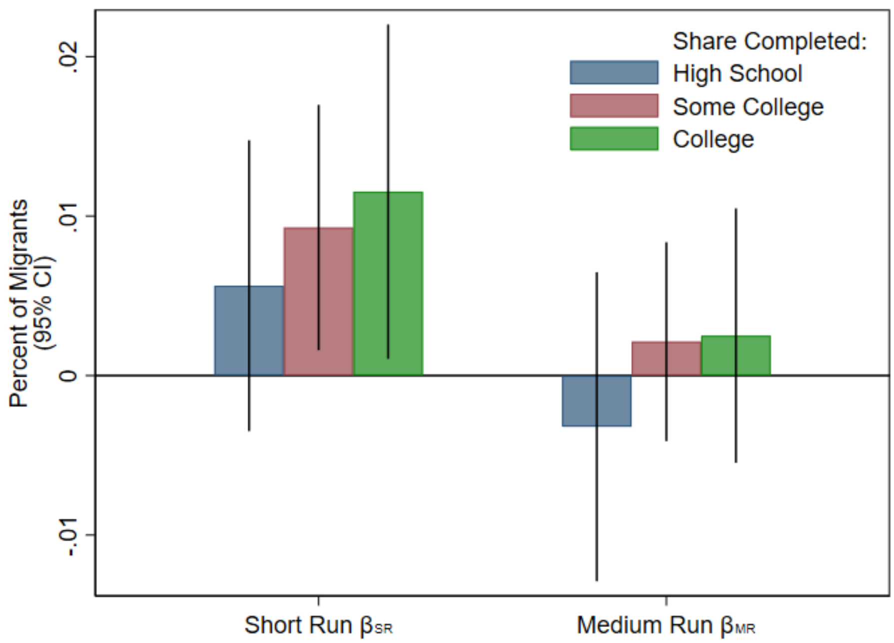

Notes: Each bar color represents coefficient estimates β_{SR} and β_{MR} from from summary specification 1.8 with different outcome variables, as noted by the legend. Unit of analysis is municipality-year. Outcome variables are the educational attainment of the stock of migrants from the 100% population census (2007, 2010, and 2015) microdata. Municipality and year-by-island-group fixed effects are included. Confidence intervals based on standard errors robust to spatial (200 km) and serial (10 year) correlation. Observations are weighted by number of migrants making up each cell.

Figure A.6: Baseline Migrant Shares are Persistent

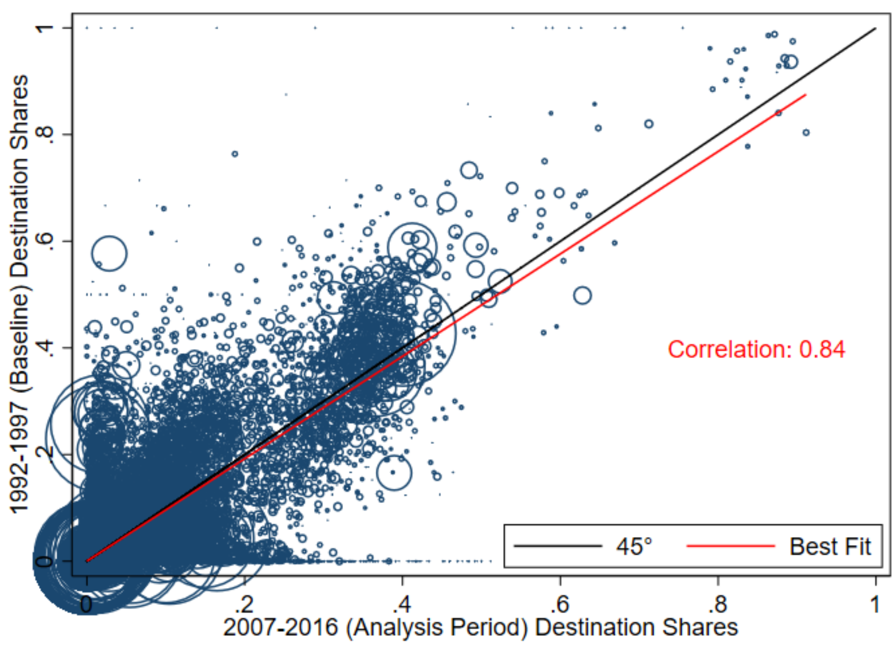

Notes: Figure plots the baseline (1992-1997) municipality-destination country migrant shares $\pi^0_{m\to d}$ and analysis period the (2007-2016) municipality-destination country migrant shares $\pi^{2007-2016}_{m\to d}$. Size of symbols reflects the total number of baseline (1992-1997) migrants from a municipality.

Figure A.7: Country Share Response to Typhoons Along the Migrant Demand Index

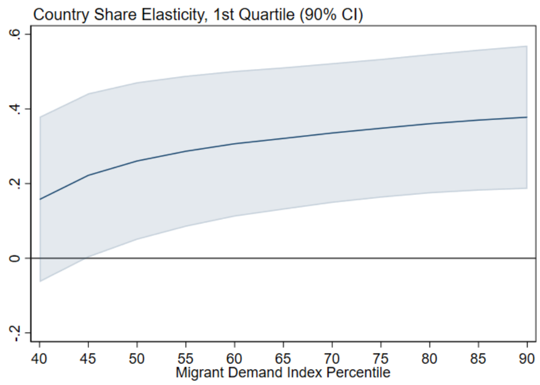

(a) 1st Country Quartile

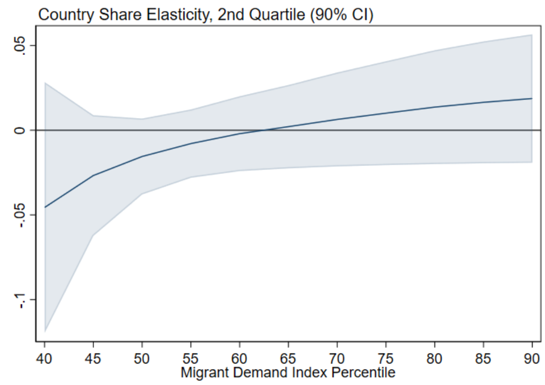

(b) 2nd Country Quartile

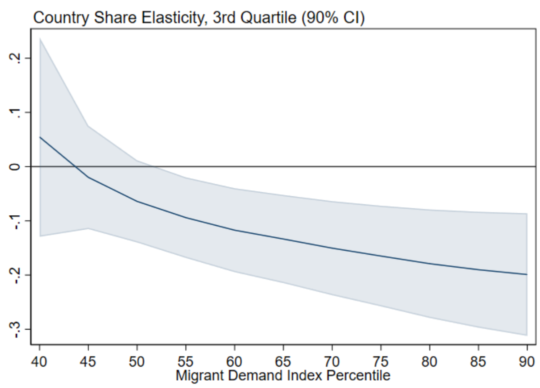

(c) 3rd Country Quartile

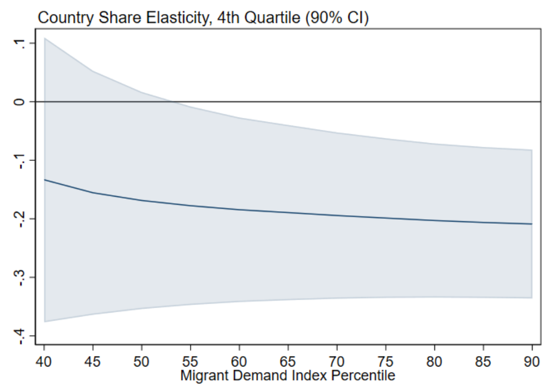

(d) 4th Country Quartile

Notes: The unit of analysis is municipality-year. Estimation includes the mean counterfactual demand index controls, corresponding to specification in columns 2 and 4 of Table 1.4. Regression results underlying the figures can be found in Appendix Table A.13. Each panel traces out the country share response heterogeneity for share of migrants going to a different quartile. Standard errors are calculated using the delta method using the variance-covariance robust at clustering at province level.

Figure A.8: Placebo: Future/Past Migrant Demand Index

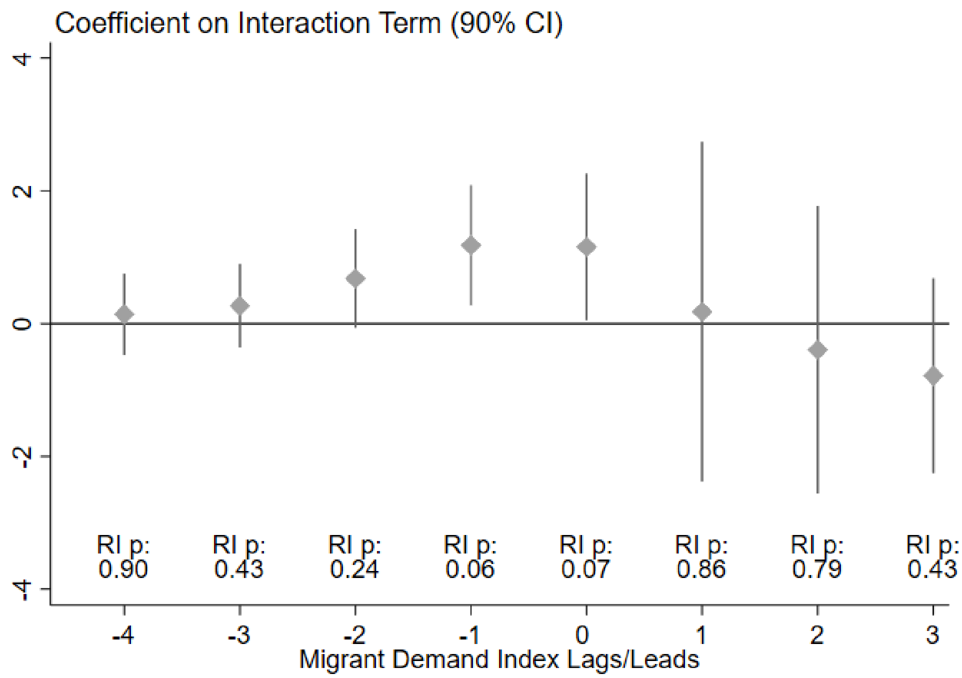

(a) Leads and Lags of Migrant Demand Index

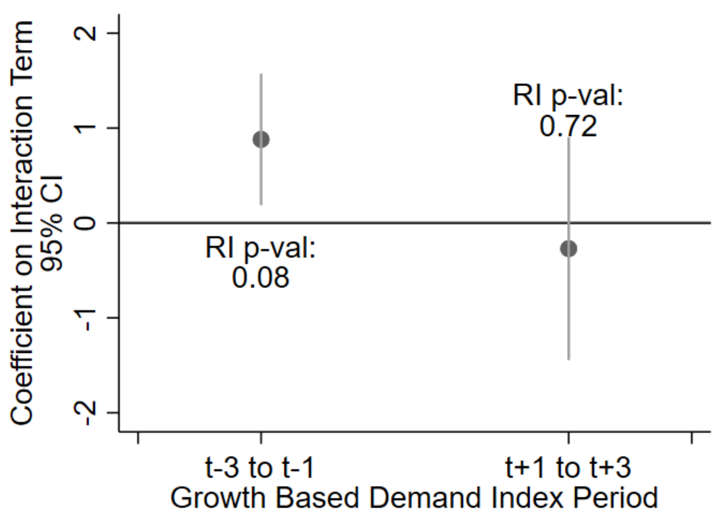

(b) GDP Growth Based Migrant Demand Index

Notes: Figures plot the interaction term of interest from estimating equation 1.12 run at the municipality level. Each estimate is from a separate regression with migration rate as the outcome variable. Regressions include the mean counterfactual demand index, corresponding to column 2 of Table 1.4. In panel (b), I recreate the demand index using the recent growth rate as opposed to levels of destination real GDP per capita. Confidence intervals are based on province clustered standard errors. Randomization inference p-values are shown on the figure.

A.2 Data and Measurement Details

A.2.1 Typhoon Exposure Measurement

I use the best-track data provided by the Joint Typhoon Warning Center (JTWC) from 2003 to 2020. This data contains meteorological information on the position, maximum sustained wind speed, radius of maximum winds, and radius of tropical storm speed winds (34 knots) at 6 hour intervals for every tropical cyclone's storm center.[1] I linearly interpolate the data to create 30 minute interval storm segment observations from the provided 6 hour interval observations.

For each 30 arc-second grid cell i, storm segment $\bar{s}$, in municipality m in year t, I calculate the predicted maximum prevailing wind speed as follows:

$$w_{i\bar{s}mt} = \mathbf{1}[mw_{\bar{s}mt} >= 34] \times \begin{cases} mw_{\bar{s}mt} & \text{if } i \text{ within radius of max winds} \\ \left[34 + (mw_{\bar{s}mt} - 34)\left(1 - \frac{d_{i\bar{s}mt}}{rad_{\bar{s}mt}}\right)^2\right] & \text{if } i \text{ between radius of max winds and radius of typhoon level wind speed} \end{cases} \tag{A.1}$$

where $w_{i\bar{s}mt}$ is the predicted wind speed, $mw_{\bar{s}mt}$ is the maximum sustained wind speed for the storm segment of interest, $d_{i\bar{s}mt}$ is the closest distance between grid cell i and maximum wind speed radius of storm segment $\bar{s}$, and $rad_{\bar{s}mt}$ is the radius of tropical storm level speed for storm segment $\bar{s}$. In words, if grid cell i is within the radius of maximum winds for storm segment $\bar{s}$, then the predicted maximum speed is the maximum wind speed of the segment. If i falls between the radius of maximum winds and the radius of tropical storm level winds, the predicted maximum prevailing speed is decaying quadratically between the two borders from the maximum wind speed to 34 knots.

Next, for each storm s, I take the maximum wind speed prevailing at each grid cell across storm segments (mw_{ismt}), subtract the threshold of 34 knots, and normalize it by the maximum wind speed observed in the data (w^{max}). The numerators and denominators are squared to account for the fact that climatologists model the impact of wind speed on structures usually with a quadratic term (Emanuel, 2011).

[1] The radius of tropical storm level speed winds is provided for each 4 quadrant (northeast, southeast, southwest, and northwest) from the center of the typhoon, which the index takes into account. For a small subset of observations, information on the the radius of tropical storm level speed winds are missing. For these observations, if possible, I interpolate the radius from surrounding observations. Otherwise, I predict the radius using a model trained on the data without the radius information missing.

$$x_{ismt} = \begin{cases} \frac{(mw_{ismt}-34)^2}{(w^{max}-34)^2} & \text{if } \geq 34 \\ 0 & \text{if } mw_{ismt} < 34 \end{cases} \tag{A.2}$$

Finally, I aggregate this storm-grid cell level information to municipality-year level index. To do so, I first use the 2000 gridded population of the world data from Socioeconomic Data and Applications Center to extract grid cell level population $N_{im,2000}$. I then take a population weighted average across storms and grid cells that are within a municipality-year and normalize this by the total population of the province:

$$T_{mt} = \frac{\sum_i N_{im,2000} \sum_s x_{ismt}}{\sum_i N_{im,2000}} \tag{A.3}$$

leaving me with the municipality-year level typhoon exposure index T_{pt} which can be viewed as an intensity-weighted typhoon exposure per capita measure. For ease of interpretation, I standardize the index to be mean 0 and SD 1 throughout the paper.

A.2.2 Migration Data and Measurement

Prior to migrating, all temporary contract migrants are required to visit the Philippine Overseas Employment Agency (POEA) to have their contract approved and to receive exit clearance. This results in POEA maintaining a database of all new temporary contract hires from the Philippines. I have access to the data set of all land-based contracts (going to a country to work as opposed to a seafarer) from 1992 to 2016. The data includes information on sex, date of birth, contract occupation, destination country, contract duration, and the wages of each migrant contract.

Migrant's origin within the Philippines is critical for my research design. Starting from 2010, the POEA data includes migrants home municipality and province. For the previous years, I rely on a matched data set between the POEA database and the database maintained by Overseas Worker Welfare Administration (OWWA), the government agency responsible for the well being of overseas workers and their families. The OWWA database includes information about migrant's home address. The matched database is created through a fuzzy matching algorithm that uses the first name, middle name, last name, date of birth, destination country, sex, and year of departure of the migrants (Theoharides, 2018a). A 95% match rate is achieved. The combined data information about migrant's home address in two periods: 1992-1997 and 2007-2016.

The 1992-1997 period as the baseline period, which I use to construct measures of baseline migration intensity and baseline province/migration-country migration shares. 2007 to 2016 is the analysis period, throughout which I observe 4 million migrant contracts, with over 90% containing origin information.

Resolving Origin Ambiguities. For 1992-2009, home address data only includes municipality (but not province) name. This creates ambiguities for municipalities with names that are repeated across provinces and affects 12% of observations in the first three years of the analysis period (the remaining 7 years don't suffer from this issue as data includes province and municipality starting in 2010). To create municipality-level variables we assign municipalities with duplicate names their population share of total population across municipalities with the same name. Where linking individuals to a specific municipality is needed (such as for individual-level analysis or calculating median wages in a municipality-year), the municipality with the highest share is assigned.

Migration Rate and Baseline Migrant Network Size. Municipality-year level migrants per capita is calculated as the sum of migrants leaving a municipality in a year divided by the population of the municipality. The populations is calculated using 2007, 2010, and 2015 censuses with linear interpolation for intermediate years. I also calculate migrant-year level (1) count of migrants, (2) migrants per 2000 capita, and (3) migrants per 2007 capita. Baseline migrant network size for municipality m is simply the sum of migrants in baseline period divided by the 1995 population:

$$NetSize_m = \frac{\sum_{t=1992}^{1997} M_{mt}}{Pop_m^{1995}}$$

where M_{mt} is the number of migrants leaving municipality m in year t and Pop_m^{1995} is the population of municipality m in 1995.

Wages. Individual migrant wages are expressed in annual terms and in real 2010 Philippine pesos. They are winsorized at 99% within each destination-occupation category (of which there are 7) cell. Where destination-occupation category cells have fewer than 100 observations, these cells aggregated and wages are winsorized at the occupation category level.

Construction of Destination Country and Occupation Quartiles. I use the baseline period data to group countries and occupations into quartiles based on their wage levels. To do so, I run a regression on log wages of contract i, to destination country d, in occupation o, in year t of the following form to jointly estimate the groupings:

$$\ln w_{idot} = \mathbf{D}_d + \mathbf{O}_o + \gamma_t + \epsilon_{iodt} \tag{A.4}$$

where $\mathbf{D}_d$ and $\mathbf{O}_o$ are the set of destination country and occupation fixed effects. I collect these estimated fixed effects and use the empirical Bayes shrinkage estimator of Morris (1983) to account for noise in the estimation leading to possible bias. I then group countries and destinations in quartiles based on the value of the fixed effects, with an equal number of occupations and countries in each quartile. I then calculate the share of migrants from a municipality in a given year that are leaving for each destination and occupation quartile simply as number of migrants going to a quartile of interest divided by total number of migrants.

A.2.3 Migrant Education

Education of Migrant Stock. I use the 100 % census data for years corresponding to my analysis period: 2007, 2010, and 2015. For each census round and municipality, I calculate the share of international migrants with high school ($\geq$ 10 years), some college ($\geq$ 11 years) and college ($\geq$ 14 years) education.

Education of New Land-Based Migrants. With the administrative data missing education data, I turn to 2011-2016 Survey on Overseas Filipinos data to calculate the educational attainment of new land-based migrants.[2]. Knowing the education, year of departure of each migrant, and geographic region of each migrant, I calculate region-year level share of land-based migrants with high school ($\geq$ 10 years), some college ($\geq$ 11 years) and college ($\geq$ 14 years) education for years 2011 to 2016.

A.2.4 Migrant Demand Index

The details of the migrant demand index is provided in Section 1.6.1. The GDP data underlying these results come from World Bank World Development Indicators (WDI). Because Taiwan is missing from the WDI, I supplement this data with Penn World Tables version 10.0 data which includes Taiwan.

[2] Distinguishing land-based vs seafarer migrants is not possible before 2011

A.2.5 Remittances

I use the Family Income and Expenditure Survey (FIES) triennial rounds between 2006 to 2018. I use income from international sources to proxy for remittances, which includes remittances but also includes pensions, retirement, and other benefits from abroad as well. Remittances are not separately reported in the data. All values are in 2010 real Philippine pesos.

A.2.6 Regression Controls

Following municipality level baseline controls are constructed from the census data:

1. Share of rural households from the 1990 Census (not available in 2000).
2. Log population from 2000 Census.
3. Share of the population with elementary school, high school, and college education from 2000 Census.

Province level baseline controls are constructed from the 2003 FIES:

1. Average per capita household income and expenditure in the province
2. Variance of per capita household income in the province

A.3 Model Derivations

This section provides the derivations underlying the model in Section 1.3. While the body provides results with a Cobb-Douglas matching function, I generalize the matching function to be any constant returns to scale function in this appendix.

A.3.1 Derivations and Proofs of Results

Model Basics. The model is in discrete time. There is a large number of homogeneous, infinitely lived, and risk neutral individuals who apply discount factor β to future utility. Individuals earn y_o per-period at home. They choose between staying at home or searching for an overseas work contract. If they search, they incur the search cost c. Finding a contract in a foreign labor market is subject to search frictions. If individuals choose to search, they match with an overseas contract with probability $q(s, v)$ which is decreasing in number of searchers s (due to congestion) and increasing in job availability v (for vacancy). Matches are realized with a CRS matching function: $m(v, s)$. The per-period probability an individual finds a contract is given by the usual matches over mass of searchers: $\frac{m(v,s)}{s} = m(\frac{v}{s}, 1) \equiv q(\theta)$ where θ is the market tightness $\frac{v}{s}$. If matched, the contract offers utility $w_t \in [\underline{W}, \overline{W}]$, drawn from the distribution $F_w(.)$. The individual can choose to accept the offer or remain home after w_t is revealed. w_t denotes the full value function associated with an offer, encoding information about its duration as well. For example, an overseas contract that offer per period utility x for two periods would have $w_t = \beta x + \beta^2 x + \beta^3 V^s$, capturing that after two periods abroad the migrant returns home and is back to searching.

Characterizing w^*. The reservation wage policy implies $w^* = V^s$. Therefore:

$$w^* = y_o - c + \beta q(\theta) \left(\int_{\underline{W}}^{w^*} w^* dF_w + \int_{w^*}^{\overline{W}} w_t dF_w \right) + \beta(1 - q(\theta))V^s$$

Subtracting $q(s, v)\beta w^*$ from both sides and rearranging terms yield:

$$w^* = \frac{y_o - c}{1 - \beta} + q(\theta) \frac{\beta}{1 - \beta} \int_{w^*}^{\overline{W}} w_t - w^* \, dF_w \tag{A.5}$$

Next, recall that in equilibrium $V^S = V^o$. Therefore, using equation A.5:

$$\frac{y_0}{1-\beta} = \frac{y_o - c}{1-\beta} + q(\theta)\frac{\beta}{1-\beta}\int_{w^*}^{\overline{W}} w_t - w^* \, dF_w$$

Note that $\int_{w^*}^{\overline{W}} w_t - w^* \, dF_w = \mathbb{E}[w_t|w_t \geq w^*](1 - F_w(w_t)) \times w^*(1 - F_w(w_t))$. I denote $\bar{F}_w(w_t) = 1 - F_w(w_t)$ as the inverse CDF. Using this expression and rearranging terms, we recover equation 1.2:

$$\beta \underbrace{q(\theta)}_{\mathbb{P}(offer)} \underbrace{\bar{F}_w(w^*)}_{\mathbb{P}(accept|offer)} \underbrace{(\mathbb{E}[w_t|w_t > w^*] - w^*)}_{\equiv \Delta(w^*)} = c$$

Finally, plugging in equation 1.2 to the above characterization of w^* implies:

$$(1-\beta)w^* = y_o - c + \beta q(\theta)\bar{F}_w(w^*)\Delta(w^*) \implies w^* = \frac{y_o}{1-\beta}$$

Average migrant wages. Given the foreign wage distribution is exogenous, $w^* = \frac{y_o}{1-\beta}$ pins down the average migrant wages as $\mathbb{E}[w_t|w_t > \frac{y_o}{1-\beta}]$.

Characterizing equilibrium migration flow m^*. Equation 1.2 implies the following expression for the total number of individuals *searching* for foreign jobs s:

$$s = v\left[q^{-1}\left(\frac{c}{\beta\Delta(w^*)\bar{F}_w(w^*)}\right)\right]^{-1} \tag{A.6}$$

Where $q^{-1}(.)$ is the inverse of the job finding rate. Total number of migrants is given by the total number of *searchers* times the job finding rate and the share of overseas offers accepted:

$$m = v\left[q^{-1}\left(\frac{c}{\beta\Delta(w^*)\bar{F}_w(w^*)}\right)\right]^{-1} q(\theta)\bar{F}_w(w^*) = v\left[q^{-1}\left(\frac{c}{\beta\Delta(w^*)\bar{F}_w(w^*)}\right)\right]^{-1}\frac{c}{\beta}\Delta(w^*)^{-1}$$

where the second equality follows from plugging in equation 1.2 for $q(\theta)$. Note that above expression implies $\frac{\partial m}{\partial v} = \left[q^{-1}\left(\frac{c}{\beta\Delta(w^*)\bar{F}_w(w^*)}\right)\right]^{-1}\frac{c}{\beta}\Delta(w^*)^{-1}$, which will be useful for pinning down equilibrium v.

To determine the vacancies, I turn to the firm problem. Firms solve:

$$\max_v pm(v,s) - v^\rho$$

where $\rho > 1$. The first order conditions imply $pm_v(v,s) = \rho v^{\rho-1}$.

Plugging in $\frac{\partial m}{\partial v} = \left[q^{-1}\left(\frac{c}{\beta\Delta(w^*)\bar{F}_w(w^*)}\right)\right]^{-1} \frac{c}{\beta}\Delta(w^*)^{-1}$, we get:

$$v = \left(\frac{p}{\rho}\right)^{\frac{1}{\rho-1}} \left(\left[q^{-1}\left(\frac{c}{\beta\Delta(w^*)\bar{F}_w(w^*)}\right)\right]^{-1} \frac{c}{\beta}\Delta(w^*)^{-1}\right)^{\frac{1}{\rho-1}} \tag{A.7}$$

Plugging in this expression to the above expression for m:

$$m = \left(\frac{p}{\rho}\right)^{\frac{1}{\rho-1}} \left(\frac{c}{\beta}\right)^{\frac{\rho}{\rho-1}} \left[q^{-1}\left(\frac{c}{\beta\Delta(w^*)\bar{F}_w(w^*)}\right)\Delta(w^*)\right]^{\frac{-\rho}{\rho-1}} \tag{A.8}$$

Taking derivatives of m and $\ln \bar{w} = \mathbb{E}[w_t | w_t > \frac{y_o}{1-\beta}]$ gives ***Result 1***:

$$\frac{\partial \ln m}{-\partial y_o} = \frac{\rho}{\rho-1}\frac{1}{1-\beta}\left[\underbrace{\frac{|\eta(\theta)^{-1}|}{\Delta(w^*)}}_{\uparrow\, search} - \underbrace{\frac{1}{\Delta(w^*)}}_{\downarrow\, \mathbb{P}(offer)} + \underbrace{\frac{f(w^*)}{\bar{F}_w(w^*)}}_{\uparrow\, \mathbb{P}(accept|offer)}\right] > 0$$

$$\frac{\partial \ln \bar{w}}{-\partial y_o} = \frac{1}{1-\beta}\frac{f(w^*)}{\bar{F}_w(w^*)}\left(\frac{w^*}{\bar{w}} - 1\right) < 0$$

where $\eta(\theta) = \frac{d\ln q(\theta)}{d\ln\theta} \in (0,1)$ is the elasticity of the job finding probability with regards to the market tightness. With a Cobb-Douglas matching function $m(v,s) = v^\alpha s^{1-\alpha}$, $q(\theta) = \theta^\alpha$, and $\eta(\theta) = \alpha$. Plugging α to above expression recovers equation 1.4 from the body of the paper.

Further taking the derivative $\frac{\partial^2 m}{-\partial y_o - \partial \rho}$ gives ***Result 2***:

$$\frac{\partial^2 \ln m}{-\partial y_o - \partial \rho} = \frac{1}{(\rho-1)^2}\frac{1}{1-\beta}\left[\frac{|\eta(\theta)^{-1}|}{\Delta(w^*)} - \frac{1}{\Delta(w^*)} + \frac{f(w^*)}{\bar{F}_w(w^*)}\right] > 0$$

Shocks to both home utility and barriers to migration. Consider the impact of a negative origin shock which has two effects. First, the shock decreases home earnings

($dy_o < 0$) which uniformly increases the returns to migration to each foreign labor market, i.e. $\frac{\partial \Delta_h}{-\partial y_o} = \frac{\partial \Delta_l}{-\partial y_o} = -1$. Second, the shock increases the fixed search costs ($dc > 0$), which is meant to capture variety of unmodeled channels a natural disaster can impede migration, such as loss of necessary documents due to damages,[3] increases in the price of migration services and assistance in response to increased demand, destruction of infrastructure making it harder to access recruitment services, or increased inability to pay the fixed cost of migration due to asset/wealth losses or disruptions to migration financing. Taking total derivative of equation A.8 for m^*:

$$d \ln m^* = \frac{\rho}{\rho - 1} \left[\underbrace{\frac{1}{1-\beta} \left[\frac{|\eta(\theta)^{-1}|}{\Delta(w^*)} - \frac{1}{\Delta(w^*)} + \frac{f(w^*)}{\bar{F}_w(w^*)} \right] (-dy_o)}_{>0} \underbrace{- \frac{1 - |\eta(\theta)|}{|\eta(\theta)|} d \ln c}_{<0} \right]$$

As expected, the change in migration is now ambiguous as the increased barriers pushes migration down while increased returns push it up (***Result 3***). The mean wages are only a function of $w^* = \frac{y_o}{1-\beta}$ and therefore are not effected by increase in migration barriers $dc > 0$.

A.3.2 Extension: Directed Search with Heterogeneous Overseas Markets

Two markets. Let s_i and m_i denote the number of *searchers* and *migrants* to overseas market $i \in \{h, l\}$, where h and l denote high and low wage overseas markets.

Assumption 1. *The high wage market has an exogenous wage distribution that is bigger than the low wage market in hazard rate ordering, i.e.* $\frac{f_w^h(w_t)}{\bar{F}_w^h(w_t)} \leq \frac{f_w^l(w_t)}{\bar{F}_w^l(w_t)}$.

Assumption 2. *The match function satisfies that* $\eta(\theta) = \frac{d \ln q(\theta)}{d \ln \theta}$ *is weakly decreasing in market tightness* θ

The first assumption introduces stochastic ordering that is stronger than first order stochastic dominance to ensure unambiguous comparative statics. The second assumption imposes enough structure on the matching function so that in more congested markets a percent change in the relative number of searchers lead to weakly lower percent change in

[3] See, for example, https://www.thenationalnews.com/world/filipinos-seek-middle-east-jobs-to-rebuild-lives-after-haiyan-1.260462 which reports in the context of the 2013 Typhoon Haiyan that "[m]any lost their passports, birth certificates and certificates of employment in the storm surge that followed the typhoon. At the Tacloban job fair, half of those applying for overseas jobs did not qualify because of a lack of documents."

the job finding probability than in less congested markets. Common matching functions in the search literature like Cobb Douglas, CES with gross complements, and urn-ball matching function satisfies this condition.

Equations A.6, A.7, and A.8 still pin down the number of searchers s_i and migrants m_i for each market $i \in \{h, l\}$. In equilibrium, the high wage overseas market attracts relatively more individuals to search for contracts, leading to lower market tightness, i.e. $\theta^h < \theta^l$. The following derivatives suffice to show whether the share of searchers and migrants going to the high wage overseas market increase in response to origin shocks dy_o:

$$d \ln \frac{s_h^*}{s_l^*} = \frac{\rho}{\rho - 1} \frac{1}{1-\beta} \left(\frac{\eta(\theta^l)^{-1}}{\Delta^l(w^*)} - \frac{\eta(\theta^h)^{-1}}{\Delta^h(w^*)} \right) dy_o > 0 \tag{A.9}$$

$$d \ln \frac{m_h^*}{m_l^*} = \frac{\rho}{\rho - 1} \frac{1}{1-\beta} \left[\frac{\eta(\theta^l)^{-1} - 1}{\Delta^l(w^*)} - \frac{\eta(\theta^h)^{-1} - 1}{\Delta^h(w^*)} + \frac{f^l(w^*)}{\bar{F}_w^l(w^*)} - \frac{f^h(w^*)}{\bar{F}_w^h(w^*)} \right] dy_o > 0 \tag{A.10}$$

which implies that the share of potential migrants searching in and migrants going to the high paying market *decreases* due to a negative home utility shock (***Result 4***). The first inequality follows from higher expected wage conditional on migration in the high wage overseas market $\Delta^h(w^*) > \Delta^l(w^*)$ (ensured by Assumption 1) and $\eta(\theta^l) < \eta(\theta^h)$ (ensured by Assumption 2). The second inequality additionally requires noting that $\eta(\theta) \in (0, 1)$ due to constant returns to scale in the matching function and $\frac{f^l(w^*)}{\bar{F}_w^l(w^*)} \geq \frac{f^h(w^*)}{\bar{F}_w^h(w^*)}$ by definition of hazard rate ordering (Assumption 1).

More than two markets. The main conclusion from the two market case is that the share of potential migrants searching in and migrants going to the high paying market *decreases* due to a negative home utility shock. This result generalizes to multiple markets.

Suppose there are many market $i \in 1, 2, ..., N$ where a lower index implies a market's wage distribution is bigger in hazard rate order that a market with higher index, i.e. $W_1 \succeq_{hr} W_2 \succeq_{hr} ... \succeq_{hr} W_N$. Equations A.6, A.7, and A.8 still pin down the number of searchers s_i and migrants m_i for each market. In response to an origin utility shock dy_0, taking derivatives parallel to A.9 and A.10 yields $d \ln \frac{m_i^*}{m_j^*} > 0$ and $d \ln \frac{s_i^*}{s_j^*} > 0$ for $i \leq j$. So, $d \ln m_1 > d \ln m_2 > ... > d \ln m_M$, which implies $d \ln \pi_1 > d \ln \pi_2 > ... > d \ln \pi_M$, where π_i is the share of migrants going to market i (same argument holds for share of searchers as well). Therefore, a negative home utility shock leads to the largest relative decline in highest paying migrant markets while leading to largest increases in lowest paying markets, i.e. a reshuffling of the migrant distribution from the high to low paying markets.

A.4 Details of Calculations

A.4.1 Calculations for Interpreting Effect Sizes in Section 1.5.6

Additional Migrants Due to Typhoons. I run the following specification analogous to main summary specification 1.8, but on the non-standardized (cardinal) values of the typhoon index T:

$$m_{mt} = \beta_{SR} T_{m,(t,t-1)} + \beta_{MR} T_{m,(t-2,t-3)} + \gamma_m + \gamma_{r(m),t} + \epsilon_{mt} \tag{A.11}$$

where m_{mt} is migration per capita from municipality m and year t. For each year and municipality, I get the predicted migration per capita due to typhoons by:

$$\hat{m}_{mt} = \hat{\beta}_{SR} T_{m,(t,t-1)} + \hat{\beta} T_{m,(t-2,t-3)} \tag{A.12}$$

Then, the total additional migrants from each municipality and year is $\hat{m}_{mt} \times pop_{mt}$ where pop is the total population of municipality in the year t. The total migration induced by typhoons in the entire analysis period is the sum over all years and municipalities:

$$\hat{M} = \sum_{t=2007}^{2016} \sum_{m} \hat{m}_{mt} \times pop_{mt} \tag{A.13}$$

Typhoon induced migration in an average year is $\frac{\hat{M}}{10}$ given there are 10 years in my analysis. This yields an average annual migrant number of 10,097.

Additional Migrant Income Due to Typhoons. The percent migrant earning response is given by percent change in migration plus the percent change in average wages of migrants, $\Delta\% me_{mt} = \Delta\% m_{mt} = +\Delta\% \bar{w}_{mt}$, where me is migrant earnings, m is migrants, and $\bar{w}$ is the average migrant wages. For each municipality and year, I calculate the percent increase in migration by:

$$\Delta\% m_{mt} = \frac{\hat{m}_{mt}}{\bar{m}_m} \tag{A.14}$$

where is the $\bar{m}_m$ is the average annual migrants from a municipality and $\hat{m}_{mt}$ is calculated from equation A.12 above.

For percent change in average migrant wages, I turn to the modified summary specification 1.8 with non-standardized (cardinal) values of the typhoon index T:

$$\bar{\ln w}_{mt} = \beta_{SR} T_{m,(t,t-1)} + \beta_{MR} T_{m,(t-2,t-3)} + \gamma_m + \gamma_{r(m),t} + \epsilon_{mt} \tag{A.15}$$

where $\bar{\ln w}_{mt}$ is the average log migrant wages in municipality m in year t. Parallel to the process for migration described above, I get the predicted changes in average migrant wages due to typhoon by:

$$\hat{\bar{\ln w}}_{mt} \equiv \Delta\%\bar{w}_{mt} = \hat{\beta}_{SR} T_{m,(t,t-1)} + \hat{\beta}_{MR} T_{m,(t-2,t-3)} \tag{A.16}$$

Total migrant earning response in a given year is then:

$$\hat{me}_{mt} = (\Delta\%m_{mt} + \Delta\%\bar{w}_{mt}) \times \bar{me}_m \times Con\bar{t}Dur_m \tag{A.17}$$

where $\bar{me}_m$ is the average total annual migrant wages in municipality m and $Con\bar{t}Dur_m$ is the average contract duration (which is around 24 months for all municipalities). The total migrant earning response throughout my analysis period is the sum of responses across years and municipalities:

$$\hat{ME} = \sum_{t=2007}^{2016} \sum_m \hat{me}_{mt} \times pop_{mt} \tag{A.18}$$

Typhoon induced migrant income in an average year is $\frac{\hat{ME}}{10}$ given there are 10 years in my analysis. This yields an average increase in total migrant earnings of 129,184,710% due to typhoons in 2010 USD. To calculate what this value would be without the migrant wage decrease due to typhoons, I set $\Delta\%\bar{w}_{mt} = 0$, which yields 157,678,510$.

A.4.2 Calculations for Share of Remittances Attributable to New Migration in Section 1.7.2

Strategy 1: Comparing levels of remittance and migrant earnings response. This calculation requires an estimate of total migrant earnings induced by typhoon exposure and total remittance increase induced by typhoons. For the migrant earnings, I follow the same strategy as in Appendix Section A.4.1 above. The only modifications is that the analysis at the province level to be consistent with the unit of analysis of the remittance responses (regression results in Appendix Table A.9).

For the remittance results, I follow an analogous strategy. I run the following specification analogous to main summary specification 1.8, but on the non-standardized (cardinal) values of the typhoon index T:

$$\ln R_{pt} = \beta_{SR} T_{p,(t,t-1)} + \beta_{MR} T_{p,(t-2,t-3)} + \gamma_p + \gamma_{r(p),t} + \epsilon_{pt} \tag{A.19}$$

where $\ln R_{pt}$ is the log remittances per capita for province p and year t. For each year and municipality, I get the predicted change in log remittance per capita due to typhoons by:

$$\hat{\ln} R_{pt} = \hat{\beta}_{SR} T_{p,(t,t-1)} + \hat{\beta}_{MR} T_{p,(t-2,t-3)} \tag{A.20}$$

Then, the total additional migrants from each municipality and year is $\hat{\ln} R_{pt} \times pop_{pt} \times \bar{R}_p$ where pop is the total population of province in the year t and $\bar{R}_p$ is the average remittance per capita of province p across years. The total remittances induced by typhoons in the entire analysis period is the sum over all years and municipalities:

$$\hat{R} = \sum_{t=2007}^{2016} \sum_{m} \hat{\ln} R_{pt} \times pop_{pt} \times \bar{R}_p \tag{A.21}$$

Typhoon induced remittances in an average year is $\frac{\hat{R}}{10}$ given there are 10 years in my analysis. This yields an average of 849,398,500$ in total, of which 458,737,800$ is due to short-run and 390,660,700$ is due to medium-run typhoon exposure.

A.5 Additional Analyses

A.5.1 Typhoon Index and Damages

To validate my constructed typhoon exposure index T, I assess the relationship between T and two outcomes: (1) province level tropical cyclone damage and casualty estimates I obtained from the Philippines National Disaster Risk Reduction and Management Council (NDRRMC) for 2003-2020 (excluding 2014) and (2) nightlight data from the Visible and Infrared Imaging Suite (VIIRS) Day Night Band (DNB) for 2012-2020.

A.5.1.1 Typhoon Index Predicts Damages and Casualties

I start with assessing whether NDRRMC damage and casualty estimates are increasing with province-year level T_{pt}. I aggregate the NDRRMC data at the province-year level. The three outcomes of interest are total number of dead, injured, and missing persons (casualties), total number of affected persons, and the total cost of the damage in real 2010 PhPs. I also run the analysis with the outcomes of interest normalized by the province population.[4] For total number of casualties and affected, I use PPML given the outcome is count data. For damages, I take a cubic root to deal with right skewness of the data. Results are highly robust to other reasonable transformations and empirical models.

Panel A of Table A.14 shows results from running a simple bivariate regression between outcome of interest and the standardized typhoon index at the province-year level. The typhoon index is highly correlated with all three outcomes of interest, with an adjusted R^2 of 47%, 19 %, and 34% for casualties, affected people, and damages respectively using the non-normalized outcomes (columns 1- 3). Panel B shows results from a panel regression of the form that includes province and year fixed effects, overall finding a similar strong association between the typhoon exposure index and government damage and casualty estimates.

A.5.1.2 Typhoon Index Predicts a Decrease in Average Nighttime Light Density

I next turn to nighttime light density. Using a processed version of VIIRS DNB that aims to remove noise, ambient light, and other background factors, I create a quarter-province level average nighttime light intensity measure by averaging the monthly NTL values of pixels falling within the borders of a Filipino province in a given quarter. Two known issues with this data is that cells with low levels of light can have negative values due to airglow contamination or noise (Samson, 2021; Uprety et al., 2019) and some monthly observations

[4]Province level GDP data to normalize the damages is unfortunately not available.

Table A.14: Typhoon Exposure Index Predicts Damage and Casualty Estimates

Panel A. Cross Sectional Regressions						
	Raw			Per Province Capita		
	Num. Casualties	Num. Affected Persons	$(\text{Damages})^{\frac{1}{3}}$	Num. Casualties	Num. Affected Persons	$(\text{Damages})^{\frac{1}{3}}$
T_{pt}	0.889***	0.484***	237.138***	0.808***	0.534***	2.929***
	(0.106)	(0.027)	(17.633)	(0.055)	(0.033)	(0.174)
Year, Prov FE	No	No	No	No	No	No
Observations	1,343	1,343	1,343	1,343	1,343	1,343
Adjusted R2	0.466	0.189	0.344	0.195	0.218	0.369
Panel B. Panel Regressions (w/ Province and Year FE)						
	Raw			Per Province Capita		
	Num. Casualties	Num. Affected Persons	$(\text{Damages})^{\frac{1}{3}}$	Num. Casualties	Num. Affected Persons	$(\text{Damages})^{\frac{1}{3}}$
T_{pt}	0.917***	0.425***	198.955***	0.899***	0.438***	2.377***
	(0.093)	(0.048)	(18.664)	(0.087)	(0.038)	(0.191)
Year, Prov FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,258	1,248	1,343	1,258	1,248	1,343
Adjusted R2			0.563			0.557

Notes: Unit of observation is province-year. Columns 1,2,4, and 5 are from a PPML model. Province clustered standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

are calculated from a very low number of cloud free day observations, potentially leading to higher variance (Skoufias et al., 2021). I replace the negative values with 0s in creating my estimate (following the approach by Skoufias et al. (2021)) and show unweighted results along with results weighted by the number of cloud free days in the province-quarter of interest, giving higher weights to presumably less noisy observations. I estimate:

$$\ln(lights)_{pt} = \alpha + \beta T_{pt} + \gamma_p + \gamma_t + \epsilon_{pt} \tag{A.22}$$

where $\ln(lights)_{pt}$ is the quarterly average nighttime light density for province p and quarter t, T_{pt} is the quarter-province typhoon exposure index (standardized to be mean 0 and standard deviation 1), and γ_p γ_t are province and quarter fixed effects. I additionally show results with a province specific trend to account for secular trends in $\ln(lights)_{pt}$ for each province, and quarter-by-island group fixed effects (consistent with my main empirical specifications) to additionally account for any island-group specific trends or shocks.

Results are presented in Table A.15. Across all specifications, I find that a one SD typhoon exposure leads to contemporaneous drop in nighttime lights by 1.7%-1.9%, further providing evidence that the typhoon exposure index captures destruction and economic damage caused by typhoons in Filipino provinces. These results are broadly in line with Strobl (2019).

Table A.15: Typhoon Exposure Index Predicts a Drop in Nightlight Intensity

	Dependent Variable: Log(Night Light Intensity)					
T_{pt}	-0.019***	-0.019***	-0.019***	-0.018***	-0.018***	-0.017***
	(0.006)	(0.005)	(0.006)	(0.006)	(0.005)	(0.006)
Prv. Trend	No	Yes	Yes	No	Yes	Yes
Year X Isl. Gr. FE	No	No	Yes	No	No	Yes
CFD Weight	No	No	No	Yes	Yes	Yes
Observations	2,765	2,765	2,765	2,765	2,765	2,765
Adjusted R2	0.921	0.934	0.941	0.931	0.948	0.952

Notes: Unit of observation is province-year. Year and province fixed effects are included. CFD stands for Cloud Free Days. Province clustered standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.10.

A.5.2 Stacked Regression with Binary Typhoon Exposure Measure

This section presents event study results for effects of typhoons on migration rates and wages using the stacked regression procedure of Cengiz et al. (2019) using a binary transformed typhoon exposure index. The goal of this analysis is to assess whether the TWFE results in the body of the paper are biased (or leads and lags in the event study graphs contaminated) due to the staggered nature of typhoon exposure and possibly heterogeneous treatment effects across provinces and time.

The analysis is done in bi-quarterly (6 month) periods. I first create a binary typhoon exposure measure using a cutoff rule, and then, for each bi-quarterly period in the analysis, construct a control group based on whether a province has had a binary typhoon exposure = 1 for 3-years before or after the bi-quarterly period of interest (referred to as a cohort).[5] Choosing the cutoff value for the binary typhoon exposure measure presents a tradeoff. Low cutoffs create a control group that is more likely to have typhoon exposure that is approximately 0, yet given the prevalence of typhoons in this setting, leads to a progressively smaller control groups as many provinces are exposed to typhoons to some degree in the 6 year event windows I construct. Smaller cutoff values also group together provinces with strong typhoon exposure with others that have progressively weaker typhoon exposures, weakening the treatment. On the other hand, high cutoffs lead to bigger control groups, yet increases the likelihood that control provinces have non-negligible typhoon exposure. Higher cutoffs also ensure that the treated cohorts have had strong typhoon exposure. Given this trade-off, I present results using two alternative cutoffs, corresponding to 95th and 90th percentile of the typhoon exposure index (conditional on non-zero exposure, these values correspond to approximately 85th and 70th percentile of the exposure index across municipality-biquarters).

For the stacked-by-event analysis, the estimating equation for the outcome y_{mtc} for municipality m, in bi-quarterly period t, for treatment cohort c:

$$y_{mtc} = \sum_{\tau=-6,\tau\neq-1}^{\tau=6} \delta_\tau \times 1[T]_{m,t,c-\tau} + \gamma_{p,c} + \gamma_{r(m),t,c} + \epsilon_{mtc} \tag{A.23}$$

where $1[T]$ is the binary treatment index and included fixed effects are also appropriately at province by treatment cohort and year by island group by treatment cohort level. Standard

[5] If treatment effects persist beyond three years, yet do not flip signs (i.e. migration response is always positive or wage response is always negative), the coefficients would be biased towards 0.

errors are clustered at the province by cohort level. I also show results from the standard TWFE specification to compare the resulting estimates to the stacked-by-event analysis estimates.

Figure A.9 presents results for the two cutoff values. There are two key takeaways. First, both sets of results suggest that the migration and wage responses are generally in line with the results in the body of the paper, though with larger magnitudes given the focus on more extreme shocks. The one departure is, for the 90th percent cutoff results, mean wages *increase* 6 bi-quarters post typhoon exposure as opposed to returning to baseline, though the same strong pattern is not present for the 95th percent cutoff results or the median wage results. Second, and more importantly, the TWFE estimates and the stacked regression estimates track each other closely. This alleviates concerns regarding whether heterogeneous effects across time or municipalities are leading to biased TWFE estimates.

Figure A.9: Stacked Regressions Results

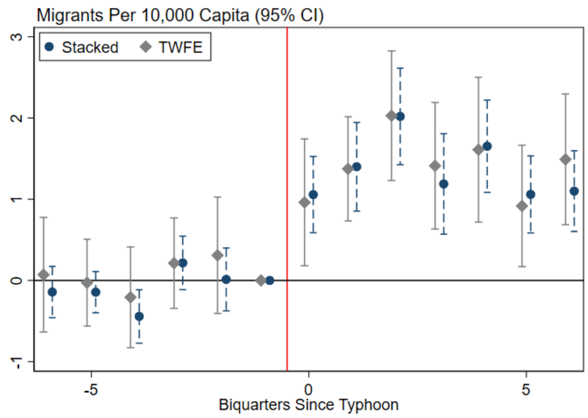

(a) Migrant Per Capita

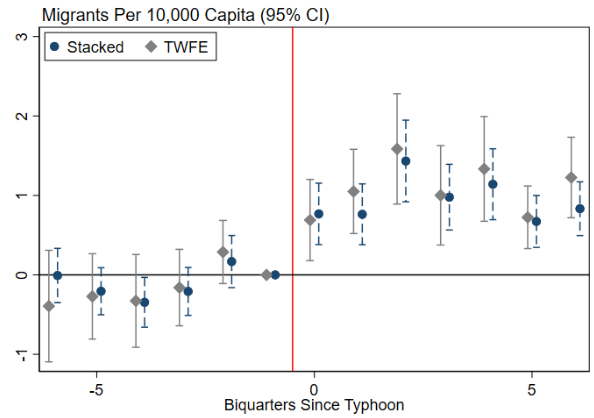

(b) Migrant Per Capita

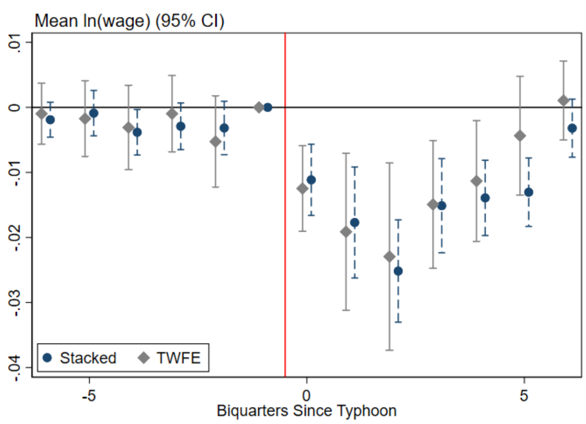

(c) Mean ln(wage)

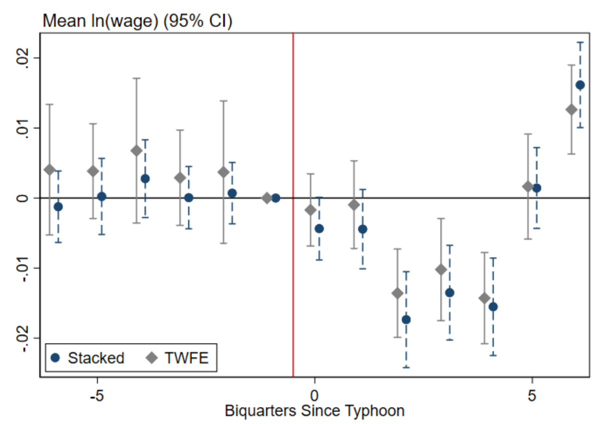

(d) Mean ln(wage)

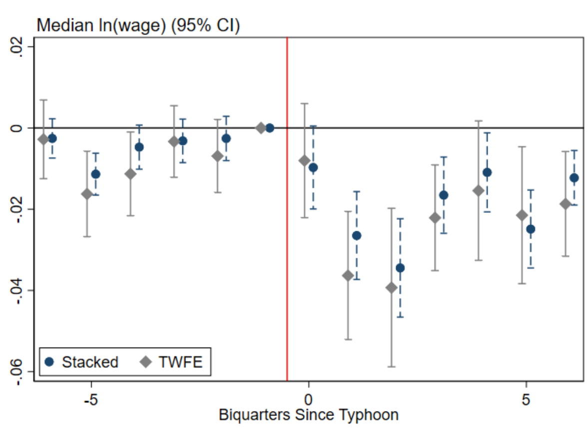

(e) Median ln(wage)

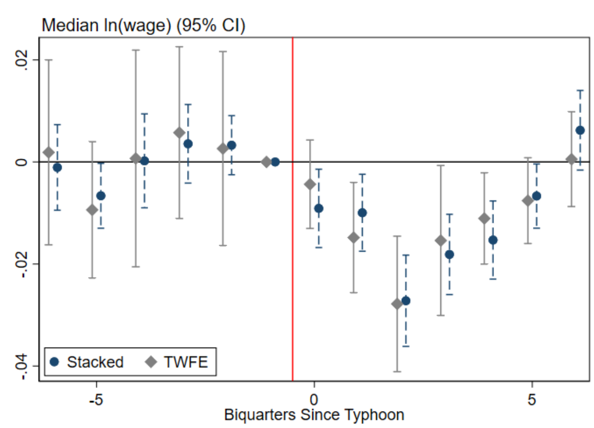

(f) Median ln(wage)

A.5.3 Variance Decomposition of Migrant Contract Wages

To document how contract wages vary with migrant characteristics present in the dataset (age and sex), destination country, and occupation, I estimate the OLS specification of the following form to undertake a variance decomposition:

$$\ln w_{idot} = \mathbf{D}_{dt} + \mathbf{O}_{ot} + \mathbf{X}_{it} + \epsilon_{iodt} \tag{A.24}$$

where w_{idot} denotes contract wages of migrant i, going to destination country d for occupation o, in year t. $\mathbf{D}_{dt}$ and $\mathbf{O}_{dt}$ are fixed effects for destination country by year and occupation by year. $\mathbf{X}_{it}$ is the full set of fixed effects for 5 year age bins interacted with sex by year. I also show results from an analogous decomposition without year interaction. The analysis covers 2007-2016.

Panels A and B of Appendix Table A.16 presents the results, with and without the year interactions. Overall, across both specifications, the destination country explains the most of the variance (33 to 36%), followed by the occupation (28 to 28.5%). The positive covariance term between occupation suggests higher paying occupation pay diferentially more in higher paying countries, though the magnitude of the covariance is relatively small. Conditional on country and occupation, age and sex by themselves explains a negligible portion of the variation (1%). 28 to 31% of the variation remains unexplained.

Given the large fraction of the variance still in the residual, I further undertake a decomposition of the form in Equation A.24 but with destination by occupation by year fixed effects. If the slope of the wage gradient of occupation are substantially different across countries, this approach can account for more of the variation than occupation and destinations individually (though this will partially be captured by the covariance term between variation explained by country and occupations in Panels A and B). Results are presented in Panel C. Occupation-destination cells now explain 75% of the variation (more than the sum of their parts and the covariance term in Panels B and C), and unexplained variation falls to 20%.

Finally, Appendix Table A.17 investigates whether contract wages differ by the origin municipality of migrants. For the set of contracts where migrant origin municipality is known, I modify Equation A.24 to include migrant origin by year dummies. Conditional on occupation, destination, and migrant demographics, inclusion of municipality makes a negligible difference in the share of variation explained, with residual share falling only by 0.2%.

Table A.16: Variance Decomposition of Migrant Wages

Panel A

	Share of Variance
Occupation	.2790469
Country	.3302533
Demographics	.0105078
Cov(Country - Occupation)	.0442084
Cov(Occupation - Demographics)	.0226242
Cov(Country - Demographics)	.0063946
Residual	.3069648

Panel B

	Share of Variance
OccupationXYear	.2851129
CountryXYear	.3604991
DemographicsXYear	.0101349
Cov(Country - Occupation)	.0414566
Cov(Occupation - Demographics)	.0218257
Cov(Country - Demographics)	.0051091
Residual	.2758617

Panel C

	Share of Variance
OccupationXCountryXYear	.7512244
DemographicsXYear	.0067424
Cov(OccupationXCountryXYear - DemographicsXYear)	.0238979
Residual	.2181353

Table A.17: Origin Municipality Does Not Explain the Wage Variation

	Share of Variance	
OccupationXCountry	.7481681	.7416341
Demographics	.0069598	.0067554
Municipality	-	.0019056
Cov(OccupationXCountry - Demographics)	.0243018	.023739
Cov(OccupationXCountry - Municipality)	-	.0066186
Cov(Municipality - Demographics)	-	.0005813
Residual	.2205704	.218766

A.5.4 Destination GDP and Migration from the Philippines

A.5.4.1 Destination GDP and Aggregate Migration from the Philippines

The migrant labor demand condition proxy described in Section 1.6 uses the destination country GDP as the country specific proxy for migrant labor demand. Therefore, it is critical to establish that destination country GDP is a meaningful proxy for migrant labor demand. This appendix investigates the relationship between aggregate migration from the Philippines and destination country GDP shocks.

I construct aggregate migration from the Philippines to individual destination countries using the administrative microdata for years 1998 to 2016. I only include countries for which there is a positive migrant flow in every year included in my analysis, which leaves me with 63 destination countries that covers over 99 % of migration out of Philippines in this period.

The estimating equation I primarily employ throughout this section is as follows:

$$\ln y_{dt} = \beta \ln GDP_{d,t-1} + \alpha_d + \alpha_t + \epsilon_{jt} \tag{A.25}$$

where y_{dt} is the outcome of interest (migration, wages, or occupation share) to country d in year t, $GDP_{d,t-1}$ is the level of real per capita GDP in country d lagged by one year, α_d and α_t are destination and year fixed effects. I refer to this model as the *panel OLS*.

Bertoli and Fernández-Huertas Moraga (2013) and Bertoli et al. (2017), shows that a reduced form estimating model of the kind A.27 implicitly restricts all destination countries to be equally substitutable, a possibly unrealistic assumption. The proposed solution is to implement the common correlated effects estimator of Pesaran (2006) which adds a set of auxiliary regressors:

$$\ln y_{dt} = \beta \ln GDP_{d,t-1} + \alpha_d + {\phi_d}' \mathbf{z_t} + \epsilon_{jt} \tag{A.26}$$

where $\mathbf{z_t} = \frac{1}{N}\left(\sum_{d=1}^{N} \ln y_{dt}, \sum_{d=1}^{N} \ln GDP_{d,t-1}\right)'$ are the cross sectional averages of the dependent and the independent variables. Inclusion of ${\phi_d}'\mathbf{z_t}$ approximates the unobserved common factors in case of cross sectional dependence, which in this context implies that the model allows for differing substitution patterns across potential destinations. I refer to this model as the *CCE*.

Migration Results. Table A.18 columns 1 and 2 shows the migration results, with log number of migrants as the outcome variable. Panel A shows results for the *panel OLS* specification with and without country specific trends. A 1% increase in lagged destination GDP leads to about 1.5% increase in out-migration across both specifications. Panel B shows that, once differential substitution patterns across destination countries are accounted for, the estimates increase to 4.7% and 3.7% with and without country trends respectively.

New Migrant Wages. Table A.18 columns 3-6 shows the wage results, with mean and median migrant wages as the outcome variable. Panels A and C provide results for raw wages, while Panels B and D shows results for wages residualized by occupation to ensure the estimates are not driven by changing occupation compositions of migrants. Across specifications, there is no evidence of a precise large positive effect, with all positive estimates insignificant at conventional levels. Without destination country trends included, wage estimates are in fact negative, though the lack of negative effects with trends suggests this may be partially picking up country level trends.

Overall, destination GDP shocks increase migration from the Philippines to the given destination, without a pronounced effect on new migrant wages. Therefore, the shocks are best interpreted as shifting the number of overseas contracts available to potential migrants.

Occupation Shares. Table A.19 presents results on whether the share of migrants leaving for high or low paying occupations change in response to GDP shocks.[6]. If GDP shocks are increasing wages alongside increasing the availability of *low-paying* occupations, the equilibrium wage effects may be masked, though occupation residualized wage results imply this is not the case. The results in Table A.19 suggests no strong consistent pattern across specifications. If anything, controlling for destination-specific linear trends, there is suggestive evidence that share of migrants going to occupations in the 2nd and 3rd wage quartiles of occupations increase due to GDP shocks, with a corresponding drop in share going to lowest paying quartile.[7]

Heterogeneity by Baseline Destination Wages. Above patterns suggest that destination GDP shocks increase migration from the Philippines to the given destination, without a pronounced effect on new migrant wages. This is consistent with an economy with binding "minimum wages" with excess migrant supply, where destination countries can increase migration in prevailing wages without increasing wages. Is this pattern heterogeneous by destination wages, where low-wage countries can have lower excess supply than high-wage countries? Appendix Figure A.10 shows results from augmenting the above OLS specifi-

[6]The occupations groupings follow the same grouping in the body of the paper, as described in Appendix Section A.2.2.

[7]This is consistent with the finding that migrant demand index primarily allows more robust migration responses and lower wage drops to typhoons through lowering the pressure to downgrade occupations.

cation with an interaction between lagged GDP and baseline average wages of destination countries $\ln \bar{w}_d^0$:

$$\ln y_{dt} = \beta_1 \ln GDP_{d,t-1} + \beta_2 \ln GDP_{d,t-1} \times \ln \bar{w}_d^0 + \alpha_d + \alpha_t + \epsilon_{jt} \tag{A.27}$$

There is no systematic relationship between wage and migration responses to destination GDP shocks, and the baseline wage levels of the destination country. This results holds for all occupations and focusing only on the lowest paying occupations.

Table A.18: Destination GDP Increases Migration, Does Not Change Wages

	ln(migrants)		mean ln(wage)		median ln(wage)	
Panel A: Panel OLS						
$\ln GDP_{d,t-1}$	1.365***	1.544**	-0.143*	0.250	-0.106	0.148
	(0.382)	(0.680)	(0.075)	(0.170)	(0.101)	(0.172)
Panel B: Panel OLS - Wage Residualized by Occupation						
$\ln GDP_{d,t-1}$	-	-	-0.173	0.134	-0.179*	0.113
	-	-	(0.107)	(0.169)	(0.100)	(0.202)
Panel C: CCE						
$\ln GDP_{d,t-1}$	4.696***	3.653**	-0.504*	0.150	-0.240	0.039
	(1.459)	(1.553)	(0.297)	(0.258)	(0.366)	(0.215)
Panel D: CCE - Wage Residualized by Occupation						
$\ln GDP_{d,t-1}$	-	-	-0.499*	-0.022	-0.253	-0.073
	-	-	(0.283)	(0.202)	(0.316)	(0.168)
Destination Ctry Trend	No	Yes	No	Yes	No	Yes
Observations	1,216	1,216	1,216	1,216	1,216	1,216

Notes: Unit of observation is destination country-year. Destination and year fixed effects are included. For wage results, observations are weighted by number of migrants going to each country, i.e. the cell size. Panel OLS corresponds to estimating equation A.27, CCE corresponds to estimating equation A.26. Destination clustered standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.10

Table A.19: Responsiveness of Migrant Occupation Shares to Destination GDP

	Share of Contracts with Occupation Quartile:							
	1st		2nd		3rd		4th	
Panel A: Panel OLS								
$\ln GDP_{d,t-1}$	-0.008	-0.321	-0.031	0.166	0.024***	0.082*	0.017	0.070
	(0.053)	(0.202)	(0.025)	(0.109)	(0.006)	(0.044)	(0.056)	(0.056)
Panel B: CCE								
$\ln GDP_{d,t-1}$	0.031	-0.085	0.014	0.081***	0.022	0.089**	-0.039	-0.003
	(0.076)	(0.123)	(0.014)	(0.017)	(0.020)	(0.041)	(0.046)	(0.051)
Destination Ctry Trend	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1,216	1,216	1,216	1,216	1,216	1,216	1,216	1,216
Adjusted R2	0.980	0.987	0.929	0.939	0.878	0.896	0.978	0.984

Notes: Unit of observation is destination country-year. Destination and year fixed effects are included. Observations are weighted by number of migrants going to each country, i.e. the cell size. Panel OLS corresponds to estimating equation A.27, CCE corresponds to estimating equation A.26. Destination clustered standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.10

Figure A.10: Aggregate Migration Response to GDP by Baseline Destination Wages

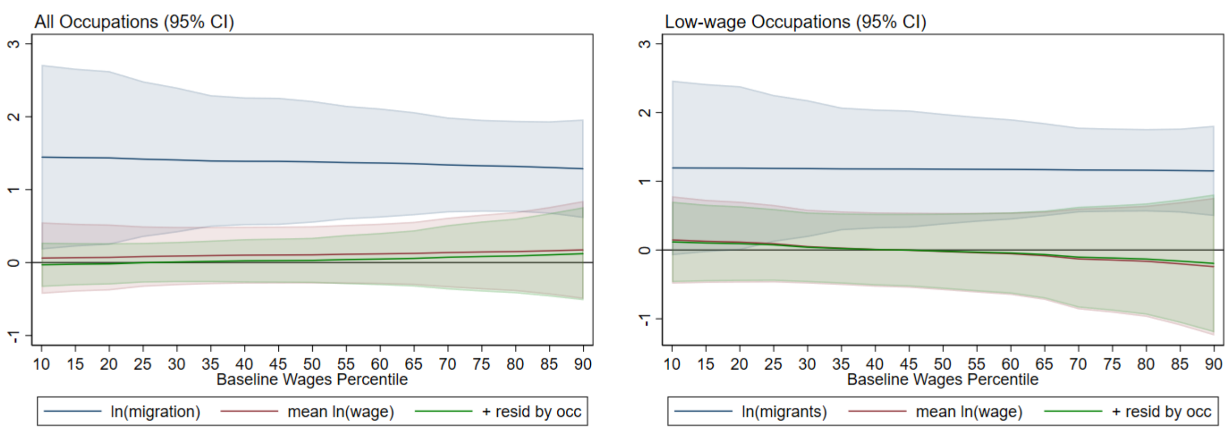

A.5.4.2 Bilateral Analysis of Migration Response to Typhoons and Destination GDP

To assess whether migration flows in response to typhoons are increasing in destination GDP, I run the following province-destination-year level bilateral regression:

$$\begin{aligned} m_{pdt} =& \beta \left(T_{p(t,t-1)} \times \ln GDP_{d,t-1}\right) + \\ & \delta \left(T_{p(t,t-1)} \times \gamma_d\right) + \gamma_{pt} + \gamma_{dt} + \gamma_{pd} + \epsilon_{pdt} \end{aligned} \tag{A.28}$$

where m_{pdt} migration flow from province p to destination country d. Due to the very high prevalence of 0s for migrant flows ($\sim$ 75% with all countries included), I deal with the right-skew of the outcome by additionally presenting results for square- and cubic-root of the migrant flows. $\ln GDP_{d,t-1}$ is the log GDP level of country d, which I lag to assuage concerns about reverse causality. $T_{p,(t,t-1)}$ is the typhoon exposure as before. The rich set of fixed effects control for the direct effects of all time-invariant province-destination characteristics (such as distance or cultural proximity), time varying province shocks (such as natural disasters themselves), and time varying destination shocks (such as any labor demand shocks or the direct effect of GDP). Finally, I control for the interaction between province typhoon exposure and destination fixed effects. This stringent control ensures that the results are not driven by migration responses differing due to any time-invariant destination characteristic. Without its inclusion a positive β could be explained, for example, by larger migration responses to countries with higher GDP overall. Inclusion of this control ensures that identifying variation for β comes from how GDP fluctuates over time within countries, as opposed to comparing across countries.[8]

Table A.20 presents the results for both all the destination countries in the data, and for the top 30 countries which account for 99% of total migration to reduce noise. Across all columns, coefficient on the interaction of interest is positive and significant, implying that migration responses to typhoons are differentially larger towards countries with high recent GDP growth. The size of the effect ranges from 10% - 30% of the mean. Therefore, GDP of destination countries are a relevant determinant of migration responses to typhoons (which bolsters the relevance of the migrant demand index). However, this result should not be interpreted as direct evidence of an increased regional migration response, as migrants may just be reallocating across destinations based on relative abundance of jobs. For the effects of migrant demand on total municipality-level migration response, see Section 1.6

[8]Results are very similar if I demean $\ln GDP_{dt}$ within each destination country as opposed to including the interaction term between typhoon exposure and destination dummy.

Table A.20: Bilateral Regression Results Between Destination Country

	All Destination Countries			Top 30 Destination Countries		
	Migrants	$(\text{Migrants})^{\frac{1}{2}}$	$(\text{Migrants})^{\frac{1}{3}}$	Migrants	$(\text{Migrants})^{\frac{1}{2}}$	$(\text{Migrants})^{\frac{1}{3}}$
$T_{p[t,t-1]} \times \text{GDP}_{d,t-1}$	14.284**	0.366***	0.134**	50.414**	1.000**	0.268**
	(6.452)	(0.135)	(0.062)	(20.972)	(0.389)	(0.118)
Observations	94,721	94,721	94,721	23,700	23,700	23,700
Adjusted R2	0.948	0.968	0.950	0.947	0.968	0.965
Mean Dep. Var.	40.411	1.920	0.926	158.413	6.589	2.848
SD Dep. Var.	386.951	6.060	1.952	761.423	10.724	3.016

Notes: Unit of observation is a province-destination-year. Typhoon exposure is at province-year level. All regressions include province-year, destination-year, and destination-province fixed effects. Province and destination country (two-way) clustered standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.10

A.5.5 Spillover Analysis

A.5.5.1 Geographic Spillovers

Economic and population ties between municipalities can lead to propagation of the economic impacts of typhoons to municipalities not directly affected by the storms. Such economic ties tend to be inversely related to

To assess this possibility, I first replicate the main migration and migrant wage analysis at the more aggregate administrative region level.[9] Using larger administrative units decreases the potential of spillovers to outside the administrative units and would further include the impacts of any within-region spillovers in the estimates. Appendix Table presents the results. Migration responses increase to 2.5 and 2.2 additional migrants per 10,000 capita to one standard deviation typhoon exposure in the short- and medium-run, corresponding to 6.7% and 5.8% of the mean (column 1). These effects are 81% and 34% larger than the municipality-level estimates, suggesting there may be positive migration spillovers. However, with log specification (column 2) the increase in estimates are more muted and not statistically significant (7% and 17%). The drop in wages in response to typhoons are also about 20% larger in the municipality-level specification as opposed to municipality-level.

Next, in appendix Table A.22, I show results from regressing migration outcomes on inverse distance weighted typhoon exposure for control municipalities (normalized to have mean 0 and standard deviation 1). In the first column, I define control municipalities as those with 0 typhoon exposure. Columns 2 and 3 defines controls less stringently by including municipalities with low direct typhoon exposure. While imprecise, Panel A shows consistently positive and economically meaningful migration impacts across the three samples, suggesting indirect exposure to typhoons through economic ties may be responding by migrating more. Panel B indicates small and imprecise migrant wage responses, especially when municipalities with low typhoon exposure are included in the control sample.

Taken together, the two analysis provides suggestive yet inconclusive evidence that indirect exposure to typhoons through economic ties (as proxied by distance) may increase migration. This spillover to control municipalities implies the main estimates presented in the body of the paper may be considered as lower bounds.

[9]Region is a more aggregate administrative unit that divides the Philippines to 17 units. I calculate the typhoon exposure analogous the municipality level exposure index.

Table A.21: Migration and Migrant Wage Responses to Typhoons - Region Level

Region Level (17 Regions)				
	(1)	(2)	(3)	(4)
			Migrant Wages	
	Migration rate	ln(migrants)	mean ln(wage)	ln(mean wage)
$T_{r,[t,t-1]}$ ($\beta_{ShortRun}$)	2.473***	0.047**	-0.012***	-0.018**
	(0.740)	(0.024)	(0.003)	(0.008)
$T_{p,[t-2,t-3]}$ ($\beta_{MediumRun}$)	2.155**	0.034**	-0.006**	-0.007
	(0.921)	(0.017)	(0.003)	(0.005)
Observations	170	170	170	170
Adjusted R2	0.923	-	0.985	0.977
Mean Dep. Var.	36.660	-	5.501	5.634

Notes: Unit of observation is a region-year. Typhoon exposure is at region-year level. All regressions include region and year-by-island-group fixed effects. Standard errors clustered at the region level in parenthesis. *** p<0.01, ** p<0.05, * p<0.10

Table A.22: Effects of Inverse Distance Weighted Typhoon Exposure

Panel A. Outcome Variable: Migration Rate			
	(1)	(2)	(3)
Sample:	$T_{m,t} = 0$	$T_{m,t} \leq 10^{th}$pctl	$T_{m,t} \leq 25^{th}$pctl
$Exposure_{m,[t,t-1]}$	1.004	1.508	1.498
	(1.812)	(1.639)	(1.654)
	[1.681]	[1.367]	[1.201]
Observations	4,310	5,278	7,129
Adjusted R2	0.814	0.832	0.846
Panel B. Outcome Variable: Mean ln(Wage)			
Sample:	$T_{m,t} = 0$	$T_{m,t} \leq 10^{th}$pctl	$T_{m,t} \leq 25^{th}$pctl
$Exposure_{m,[t,t-1]}$	0.006	0.001	0.001
	(0.009)	(0.010)	(0.005)
	[0.011]	[0.011]	[0.006]
Observations	4,200	5,148	6,986
Adjusted R2	0.867	0.896	0.898

Notes: Unit of observation is a municipality-year. All regressions include municipality fixed effects, year-by-island-group fixed effects, and the direct effect of typhoon exposure for columns 2 and 3. Exposure is defined as the inverse distance weighted average of the typhoon exposure of all other municipalities. Exposure is normalized to have mean 0 and standard deviation 1 Standard errors clustered at the region level in parenthesis. *** p<0.01, ** p<0.05, * p<0.10

A.5.5.2 Effects of Typhoons to Provinces with Similar Migrant Destinations and Recruitment Agencies

If migrants across municipalities are competing for same overseas contracts, typhoon exposure in one part of the country can increase competition for the contracts in other parts, leading to possibly negative spillovers. Consider the extreme case that total number of overseas contracts in the Philippines each year is fixed. Then, typhoons could reallocate contracts to typhoon-struck regions due to increased search intensity, yet not increase aggregate migration. On the other hand, if each province is an "island" with its own supply of possible overseas contracts, such spillovers would not occur.

Assessing the empirical relevance of such a spillover is challenging. We would need variation on which provinces are differentially competing with each other for overseas contracts. I aim to make progress by (1) leveraging the persistence of province-destination ties and (2) availability of recruitment agency data in the baseline period. Specifically, I calculate how "similar" the baseline destination networks or recruitment agency networks of any two provinces. Then, I create a province level exposure measure that weights the typhoon exposure of all other provinces by how "similar" their destination or recruitment agencies are. I further weigh each provinces contribution by the total number of migrants leaving the municipality in 1992-1997, as provinces with more migrants would lead to more competition. I then regress migration outcomes of a municipality on these exposure measures. I discuss the specifics of the calculations below.

Baseline Destination Network Similarity Weighted Typhoon Exposure. I start by calculating the baseline period share of migrants from each province p going to each destination country d, $\pi^d_{p\to d}$. This provides me with the vector $\pi^{\mathbf{d}}_{\mathbf{p}} = [\pi^d_{p\to 1}, ..., \pi^d_{p\to D}]$, which denotes the baseline destination network of each province. I calculate the pairwise destination network similarity across all provinces using cosine similarity:

$$sim^d_{ij} = \frac{\pi^{\mathbf{d}}_{\mathbf{i}} \cdot \pi^{\mathbf{d}}_{\mathbf{j}}}{\|\pi^{\mathbf{d}}_{\mathbf{i}}\|\|\pi^{\mathbf{d}}_{\mathbf{j}}\|}$$

I then define a typhoon exposure measure for each province-year that is the weighted average of the typhoon exposure of all other provinces weighted by the cosine similarity and the baseline number of migrants (M^0_i) from the province:

$$Exposure^d_{p,[t,t-1]} = \sum_{i\neq p} \frac{sim^d_{pi}}{\sum_{j\neq p} sim^d_{pj}} T_{i,[t,t-1]} M^0_i$$

where the weighting by total number of baseline migrants M^0_i aims to capture that provinces with more migrants would lead to more congestion in the overseas labor markets.

Baseline Recruitment Agency Similarity Weighted Typhoon Exposure. I calculate the recruitment agency similarity analogous to baseline destination country share. I first calculate the baseline period share of migrants from each province p using recruitment agency r, $\pi^0_{p\to r}$. This provides me with the vector $\pi^{\mathbf{r}}_{\mathbf{p}} = [\pi^r_{p\to 1}, ..., \pi^r_{p\to D}]$, which denotes the baseline destination network of each province. I again calculate the pairwise destination network similarity across all provinces using cosine similarity:

$$sim^r_{ij} = \frac{\pi^{\mathbf{d}}_{\mathbf{i}} \cdot \pi^{\mathbf{r}}_{\mathbf{j}}}{\|\pi^{\mathbf{r}}_{\mathbf{i}}\| \|\pi^{\mathbf{d}}_{\mathbf{j}}\|}$$

I then define a typhoon exposure measure for each province-year that is the weighted average of the typhoon exposure of all other provinces weighted by the cosine similarity and the baseline number of migrants (M^0_i) from the province:

$$Exposure^r_{p,[t,t-1]} = \sum_{i\neq p} \frac{sim^r_{pi}}{\sum_{j\neq p} sim^r_{pj}} T_{i,[t,t-1]} M^0_i$$

where, again, the weighting by total number of baseline migrants M^0_i aims to capture that provinces with more migrants would lead to more congestion in the overseas labor markets.

Controls. One concern with the exposure measures above is that they can be correlated with the direct typhoon shock to the province of interest itself and destination/recruitment agency similarity is correlated geographic distance and therefore economic ties. Therefore, exposure may be correlated with direct impacts of a typhoon to a province or its economic impacts through effecting provinces with economic ties, leading to an omitted variable issue. To deal with this, I include controls for direct typhoon exposure and additionally include three controls:

1. Inverse distance weighted typhoon exposure of other provinces.

2. Average typhoon exposure of neighboring provinces.

3. Baseline migrant share weighted average of typhoon exposure of other provinces.

Results. Appendix Table A.23 presents the results. Neither exposure measure predicts lower migration rate from a province, with all point estimates positive yet imprecise. Therefore, this analysis does not yield any evidence of negative spillovers in terms of migrant flows. Interestingly, I find small but negative effects on average wages with both exposure measures in the short run. Combined with positive estimates for migration, this may suggest that recruitment agencies respond to typhoon driven shift in migrant supply by securing additional (possibly lower paying) overseas contracts. Of course, if provinces with similar recruitment agency/destination structure have close economic ties, I may still be capturing economic disruptions at home if the included controls do not perfectly control for such ties. Overall, while not conclusive, results point against negative spillovers due to increased congestion in overseas contract markets.

Table A.23: Spillovers: Effects of Exposure As Measured by Baseline Destination Country/Recruitment Agency Network Similarity

Panel A. Exposure: Baseline Destination Network Similarity									
	Migration Rate			Mean ln(salary)			Median ln(salary)		
$Exposure_{p,[t,t-1]}$	0.703**	0.500	0.438	-0.002	-0.002*	-0.003*	-0.002	-0.002	-0.002
	(0.300)	(0.307)	(0.296)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
$Exposure_{p,[t-2,t-3]}$	0.329	0.206	0.206	-0.001	-0.001	-0.000	0.002	0.002	0.002
	(0.352)	(0.353)	(0.360)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
$T_{p,[t,t-1]}$		1.120**	0.989*		-0.011***	-0.010**		-0.015***	-0.014**
		(0.505)	(0.563)		(0.003)	(0.004)		(0.004)	(0.007)
$T_{p,[t-2,t-3]}$		1.108	1.130		-0.006***	-0.006***		-0.013**	-0.013**
		(0.732)	(0.720)		(0.002)	(0.002)		(0.006)	(0.006)
Network Controls	No	No	Yes	No	No	Yes	No	No	Yes
Observations	790	790	790	790	790	790	790	790	790
Adjusted R2	0.936	0.937	0.937	0.970	0.971	0.972	0.920	0.923	0.923
Panel B. Exposure: Baseline Recruitment Agency Network Similarity									
	Migration Rate			Mean ln(salary)			Median ln(salary)		
$Exposure_{p,[t,t-1]}$	0.536*	0.370	0.294	-0.002	-0.002*	-0.003**	-0.001	-0.002	-0.002
	(0.279)	(0.270)	(0.250)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
$Exposure_{p,[t-2,t-3]}$	0.257	0.140	0.146	-0.001	-0.001	-0.000	0.001	0.001	0.001
	(0.349)	(0.342)	(0.350)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
$T_{p,[t,t-1]}$		1.219**	1.055*		-0.011***	-0.009**		-0.015***	-0.014**
		(0.499)	(0.560)		(0.003)	(0.004)		(0.004)	(0.007)
$T_{p,[t-2,t-3]}$		1.215	1.231*		-0.006***	-0.006***		-0.013**	-0.013**
		(0.735)	(0.723)		(0.002)	(0.002)		(0.006)	(0.006)
Network Controls	No	No	Yes	No	No	Yes	No	No	Yes
Observations	790	790	790	790	790	790	790	790	790
Adjusted R2	0.936	0.937	0.937	0.970	0.972	0.973	0.919	0.923	0.923

Notes: Unit of observation is a province-year. All regressions include province fixed effects and year-by-island-group fixed effects. Standard errors clustered at the province level in parenthesis. *** p<0.01, ** p<0.05, * p<0.10

A.5.6 Alternative Randomization Inference Procedures

To check the robustness of my conclusions to the assumed DGP underlying the randomization inference procedure, Appendix Table A.12 presents randomization inference p-values for the interaction terms of interest from alternative assumption about the DGP. For each DGP, I present results for when the error term is sampled from the empirical distribution of the error terms in each year ("empirical distribution") and when the error term is distributed from a random distribution matching the variance and mean (which is mechanically 0) of the error terms each year ("normal distribution"). The DGPs are as follows:

- Column (1) is the baseline case where DGP is assumed to be:

$$\ln(g_{ct}) = \theta \ln(g_{ct-1}) + \delta_t + \epsilon_{ct}$$

- Column (2) includes an AR2 term:

$$\ln(g_{ct}) = \theta_1 \ln(g_{ct-1}) + \theta_2 \ln(g_{ct-2}) + \delta_t + \epsilon_{ct}$$

- Column (3) includes an a destination country fixed effect capturing the average growth rate of the country in the period:

$$\ln(g_{ct}) = \theta_1 \ln(g_{ct-1}) + \delta_c + \delta_t + \epsilon_{ct}$$

- Column (4) includes a gulf country by year control allowing for gulf countries to have correlated shocks:

$$\ln(g_{ct}) = \theta_1 \ln(g_{ct-1}) + \gamma \mathbf{1}[GCC_c] \times \delta_t + \delta_c + \delta_t + \epsilon_{ct}$$

- Column (5) includes all four additions:

$$\ln(g_{ct}) = \theta_1 \ln(g_{ct-1}) + \theta_2 \ln(g_{ct-2}) + \delta_c + \gamma \mathbf{1}[GCC_c] \times \delta_t + \delta_t + \epsilon_{ct}$$

A.5.7 Response Heterogeneity by Historical Migration

Established migrant networks can help facilitate future migration by reducing migration costs (Munshi, 2003a; Beine et al., 2011). For example, previous migrants tend to be important sources of information. The 2015 and 2016 KNOMAD/ILO surveys finds that about half the Filipino labor migrants in the sample learned about their current overseas job from relatives or friends. Further, high-levels of past migration indicate more recruitment agency activity in a given area, making it easier for potential migrants to search, apply and leave for overseas contracts. The intensity of past migration can therefore allow for a larger migration response through lower migration costs.[10] Can lower search and information costs also dampen the need to downgrade occupations and countries? Or does the possibly larger migration response come at the cost of further decreasing average wages due to aggregate limitations on job availability? To explore these questions, I check whether migration responses are heterogeneous by baseline migration intensity of a municipality. I use the baseline period data to create a municipality's migrant network size measured as total migrants from 1992-1997 per 1995 population ($\frac{\sum_{t=1992}^{1997} M_{mt}}{Pop_m^{1995}}$).

Figure A.11 presents the short-run migration rate and migrant wage results for each quintile of migrant network size.[11] Migration responses to typhoons are increasing in past migration. This is consistent with the findings of Mahajan and Yang (2020a) and Winter (2020), though in a context where the specific mechanisms they document (family migration enabled through existing family members) can't be operational. However, the drop in migrant cohort wages are also concentrated in the high baseline network municipalities. Any improved search conditions at home due to past networks does not seem to allow the municipality to increase migration without a proportional drop in wages. This is in contrast with the seminal results of Munshi (2003a) who shows bigger migrant network in US destinations lead to better employment and wage outcomes for Mexican migrants.

[10] However, a bigger existing migrant network can also decrease the incentives to additional migration by allowing for more robust insurance through remittances.

[11] Appendix Figure A.12 shows robustness to an alternative migrant network measure using 1995 census data.

Figure A.11: Migration and Wage Response to Typhoons are Increasing in Baseline Migrant Network Size

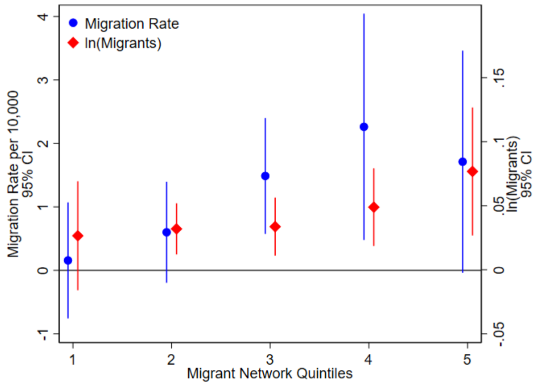

(a) Migration

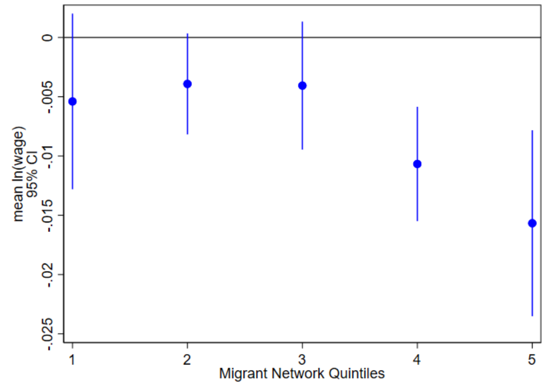

(b) Migrant Wages

Notes: Coefficients estimates for the interaction between short run typhoon exposure $T_{m,(t,t-1)}$ and dummies for baseline network size quintiles are plotted. Unit of analysis is municipality-year for both panels. Municipality and year-by-island-group fixed effects are included. Confidence intervals based on province clustered standard errors. Log migrant regression is estimated using PPML in panel (a). Observations are weighted by number of migrants making up each cell in panel (b).

Figure A.12: Robustness: Baseline Migrant Network Size Results Using Alternative Measurement from 1995 Census

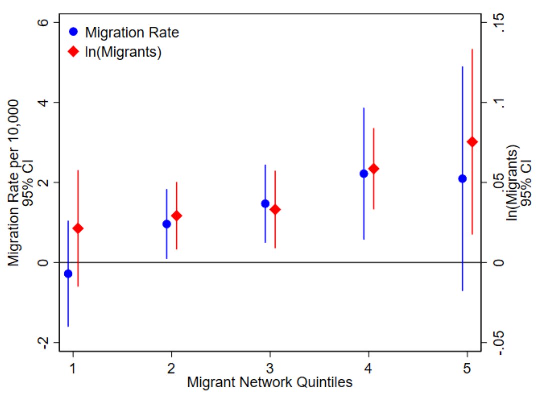

(a) Migration

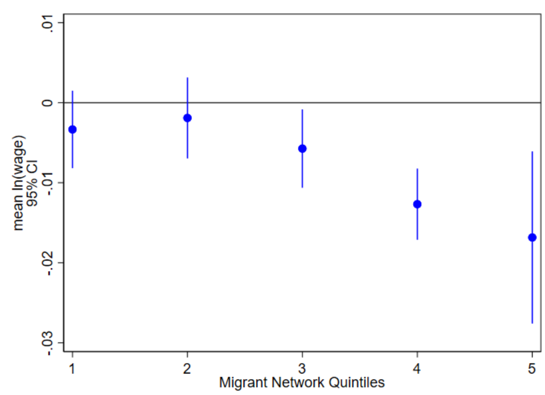

(b) Migrant Wages

Notes: Coefficients estimates for the interaction between short run typhoon exposure $T_{m,(t,t-1)}$ and dummies for baseline network size quintiles are plotted. Unit of analysis is municipality-year for both panels. Municipality and year-by-island-group fixed effects are included. Confidence intervals based on province clustered standard errors. Log migrant regression is estimated using PPML in panel (a). Observations are weighted by number of migrants making up each cell in panel (b).

APPENDIX B

Appendix to Chapter II

B.1 Data Appendix

B.1.1 Migration Data

Calculation of migrant income per capita of each Philippine province in every overseas destination requires unusual data. We obtained two administrative datasets from Philippine government agencies. The Philippine Overseas Employment Administration's (POEA) migrant contract database contains name, date of birth, sex, marital status, occupation, destination country, employer, recruitment agency, salary, contract duration, and date deployed. The database of the Overseas Worker Welfare Administration (OWWA) includes migrants' name, date of birth, sex, destination country, date deployed, and home address in the Philippines.

To create a dataset that includes migrant wages, destination, and province of origin, we combine the datasets from POEA and OWWA using fuzzy matching techniques for the years 1992-1997 and 2007-2009. We match the POEA and OWWA data using first name, middle name, last name, date of birth, destination country, sex, and year of departure. We achieve a match rate of 95%. Starting in 2010, data from POEA included wages, destination, and province of origin, so our data from 2010-2015 is from POEA only and does not require matching. Several of the immediate post-shock (post-1997) years have relatively high rates of missing data on migrant origin address. We therefore focus on the years 2007-2015 which have low rates of missing address data, and which also span the 2007, 2010, and 2015 Philippine Censuses.[1] All wages are expressed in thousands of real 2010 Philippine pesos.

[1] In the 1992-2009 contract data, the home address variable in the OWWA data includes municipality, but not province. Out of 1630 municipalities in the Philippines, 332 have names that are duplicated in another province. This accounts for between 10 and 19% of migration episodes depending on the year. Thus, to calculate province-level variables, we assign municipalities with such duplicate names their population share of the total wages across municipalities with the same name. For the 2010-2015 data, municipality and province are reported for each contract.

We winsorize the wages at 99% within each destination-occupation category cell.[2]

We use the 1995 contract data to construct the shift-share variable $Shiftshare_o$. First, we calculate province-level migrant income per capita ($MigInc_{o0}$) in 1995. We calculate province total migrant income by multiplying average migrant income for a province's migrants in 1995 (from the POEA/OWWA contract data) by the number of migrants in a given province (from the 1995 Census). We then divide by 1995 province population, obtaining migrant income per capita. We use an analogous calculation for migrant income per capita in 1994, 2009, 2012, and 2015 (corresponding to triennial FIES years). For each year, we calculate average migrant income from the POEA/OWWA data.[3] We then multiply by the total number of migrants in the 1995 Census (for 1994 migrant income per capita), 2010 Census (for 2010 and 2012 migrant income per capita) or in the 2015 Census (for 2015 migrant income per capita).

Second, we use the contract data to construct $Rshock_o$, the weighted average exchange rate shock of province o's migrants. Weights are pre-shock share of migrant income from destination d. For each province o, we calculate these weights directly from the contract data, as the share of total province-level migrant annual income from each destination country in 1995 ($\frac{\omega_{do0}}{\sum_d \omega_{do0}}$). We then multiply each exchange rate change $\tilde{\Delta} R_{d0}$ by the corresponding province-o-specific weights to obtain $Rshock_o$.

A small minority of contracts have missing data on municipality in the OWWA data (14.5% in 1995). A concern is that the exchange rate shock might be correlated with the propensity to be missing municipality data in the pre-period, and thus introduce some chance correlation with province or destination characteristics into $Shiftshare_o$. To test this, we regress the exchange rate shock on the share of destination observations with a missing province on the exchange rate shock, weighting by Borusyak et al. (2022a) shares. The regression specification is the same as in Appendix Table B.1. The coefficient on the share missing is very small in magnitude and not statistically significantly different from zero. A one-standard-deviation increase in the share of contracts missing province data is associated with a 0.007 increase in the exchange rate shock (which has a mean of 0.406 and a standard deviation of 0.138). The regression provides no indication that the propensity for migrant worker contracts for a given migration destination to have missing Philippine location data in the pre-period is correlated with that destination's exchange rate shock.

[2]When destination-occupation cells have fewer than 100 observations, we aggregate these cells and winsorize at the occupation level.

[3]For these years, we use the migrant wages from the previous three years of contract data to calculate average income per migrant. For example, 2009 migrant income per capita uses the average of income reported in contracts in 2007, 2008, and 2009. Migrant contracts have an average contract length of 24 months, so the average wages of the stock of migrants in 2009 would reflect the average wages of migrants departing in 2009 as well as previous years.

B.1.2 Domestic Income and Expenditure

All outcomes in money units in this paper (e.g., income and expenditure) are in 2010 real Philippine pesos (PhP; 17.8 PhP per PPP US$ in 2010).

Data on household income and expenditure are from triennial rounds of the Philippine Family Income and Expenditure Survey (1985, 1988, 1991, 1994, 1997, 2000, 2003, 2006, 2009, 2012, 2015, and 2018). The FIES provides the Philippine government's official income and expenditure statistics. It includes detailed household income and expenditure items. Domestic income and expenditure (as in Table 2.3), are the aggregation of these detailed items. Domestic income is calculated as total household income minus income from international sources, transfers from domestic sources, and gifts from other households. Income from international sources includes migrant remittances, but also includes pensions, retirement, workmen's compensation, and other benefits; cash gifts, support, relief, etc. from abroad; and dividends from investments abroad. Migrant remittances are not explicitly reported in the data.

We calculate global income by adding migrant income from the POEA/OWWA data and domestic income from the FIES. To analyze global income's domestic and migrant components, we focus on a subset of time periods when both domestic and migrant income data are available. This allows us to examine one pre-shock year and three post-shock years in analyses of global income. For domestic income from the FIES, the pre-shock year is the 1994 FIES round, and the post-shock years are 2009, 2012, and 2015 FIES rounds.

B.1.3 Census Data

We created a panel of schooling outcomes using the 1990, 1995, 2000, 2007, 2010, and 2015 Philippine Census of Population. In each census round, we calculate the provincial share of individuals with primary (6 or more years of schooling), high school (10 or more years), and college education (14 or more years) for the full population (aged 20-64) as well as for international migrant workers.

B.1.4 Labor Force Survey Data

The FIES, which we use for our main income and expenditure outcomes, is implemented as a rider every three years to the government's quarterly Labor Force Surveys (LFS). We use the merged LFS and FIES data to calculate domestic income per capita for skilled and unskilled households (used in the model-based quantification, Appendix Section B.2). The LFS indicates the education level and the employment status of each member of the household. We define a household as "skilled" if any of the employed members have a college education or above. We then calculate domestic income per capita for skilled and unskilled households using the FIES.

B.1.5 Data for Quantifying Contribution of the Education Channel

We create a database at the origin-destination-skill group-by-year level from our raw data in order to carry out the model-based quantifications. We use the 1990 Census to construct the baseline probability of migration by skill-group (shares of working-age population who migrated, by skill group). In addition, we use the POEA/OWWA data to construct migrant income for each origin-destination pair, by skill group and year. We use the post-shock period to determine the returns to skill using these incomes. We exclude origin-destination-skill-time observations where there were no flows. We winsorize the salary data at the 99th percentile.

B.1.6 Regression Controls

B.1.6.1 Destination-Level Controls

Destination-level controls are aggregated to the province level by taking weighted averages of destination-level variables for each province, weighted by baseline migrant earnings from each destination, following Borusyak et al. (2022a). To construct baseline GDP per capita, we used 1995 values in current US dollars from World Development Indicators.[4] The baseline destination contract variables are the following four variables from the 1995 POEA/OWWA data: (1) average 1995 salary (in real 2010 Philippine pesos) for each destination's contracts, (2) percent of 1995 contracts in professional occupations, (3) percent of 1995 contracts in production occupations, and (4) percent of all 1995 contracts for Philippine international migrant workers going to the destination.

[4]For the following small set of destinations this variable was not available in the WDI. For Taiwan, we used 1995 GDP per capita values from Taiwan's national statistics https://eng.stat.gov.tw/ct.asp?xItem=37408&CtNode=5347&mp=5. For Guam, Midway Island, and Northern Mariana Islands, we used US baseline values as they are US territories. For British Overseas Territories Cayman Islands and Diego Garcia we use UK baseline values. For Netherlands Antilles, we used Netherlands baseline values. For Palau, we use the closest available year of 2000 GDP per capita. Finally, Netherlands and Myanmar only had 1995 GDP per capita in 2010 US$ and had 1999 GDP per capita in current US$ (the closest year to 1995). We used the following estimate: $gdppc_{1995}^{currentUS\$} = gdppc_{1995}^{2010US\$} \times \frac{gdppc_{1999}^{currentUS\$}}{gdppc_{1999}^{2010US\$}}$

B.1.6.2 Province-Level Controls

Baseline share of rural households is from the 1990 census. Baseline asset index is from the 1990 census. This is the first principal component of household-level indicators for ownership of a set of durable goods, utilities access, housing quality, and land and home ownership. We then take the mean of this household-level index within each province. Baseline domestic income and expenditure per capita are the average of domestic income per capita and expenditure for 1988, 1991, and 1994, calculated from FIES microdata. Baseline sector shares are shares of employed individuals in primary, industrial, service, and financial/business services sectors, calculated from the 1990 census.

B.1.7 Exports and Foreign Direct Investment

In Section 2.5.3, we examine potential other mechanisms for our causal effects: manufactured exports and foreign direct investment (FDI).

Data on manufacturing firm exports are from a set of firm sample surveys of the Philippine Statistics Authority: the Annual Survey of Establishments (1994, 1996, 1997, 1998), Annual Survey of Philippine Business and Industry (2008, 2009, 2010, 2013, 2014, 2015), and Census of Philippine Business and Industry (1999, 2006, 2012). We obtain data for province-year observations that had three or more manufacturing establishments in the sample.[5] We sum exports across firms to the province-year level, then divide by province population to obtain per capita figures. Summed exports within province-year cells account for survey sampling weights when available (2000 and after). (Results are robust to using unweighted sums for all years.) We winsorize province-year observations at 99%.

FDI data for 1996-2002 are available from the PSA's Foreign Investment Reports, which provide the breakdown of total approved foreign investments by origin country. FDI data for 2003 and after are from the PSA's OpenStat platform. Data on FDI is broken down at the country level for major investors. FDI coming from other countries are not broken down by country and are assumed to be zero in the analysis.[6]

[5] Data are not released for province-year cells with fewer than three firms, for confidentiality reasons. We impute zeros for these province-year observations.

[6] The average share of yearly FDI not broken down by country is 6.9%.

B.2 Model-Based Quantification: Full Elaboration of Model

We present a theoretical framework relating migrant exchange rate shocks to domestic and migrant income. We use this framework to derive our empirical specification and interpret our findings. We build on recent gravity models (Bryan and Morten, 2019; Tombe and Zhu, 2019) which adapt Eaton and Kortum (2002) to model migration. We endogenize skill investments, and allow for skill-dependent migration and income, to further deepen our understanding of mechanisms and magnitudes. Full derivations of the model equations are in Appendix B.3.

We start by introducing the migration decision, and how the migrant income shock helps us derive the empirical independent variable of interest: the shift-share we use for estimation. Then we study educational investments in the theoretical model, and we estimate our gravity equation to quantify the elasticity of migrant flows with respect to destination wages. With these estimates at hand, we evaluate the effects of the exchange rate shock on origin province migrant flows, migrant income, and domestic income in our model and quantify the importance of the education channel.

B.2.1 Migration Decisions

An individual i's earnings vary across origin province o, destination country d, skill level s, and time t. They depend on destination-specific wage profiles w_{dst} (wages in destination differing by skill) and exchange rates R_{dt}. Additionally, ϵ_{dot} is any unobservable factor that makes migrants from origin o more productive in destination d. Overseas wages w_{dst} and unobservable component ϵ_{dot} are in destination-d currency units. Exchange rates R_{dt} are in Philippine pesos (PhP) per destination-d currency unit. We denote $w_{dost} \equiv w_{dst}\epsilon_{dot}$ as the wage profiles of workers from o in destination d.

Individuals have destination-specific preference draws q_{id}. Workers lose a fraction of their earnings to migration cost $0 \leq \tau_{dot} \leq 1$. Indirect utility from destination choice is:

$$V_{idost} = w_{dst}\epsilon_{dot}R_{dt}(1 - \tau_{dot})q_{id} \equiv w_{dost}R_{dt}(1 - \tau_{dot})q_{id} \tag{B.1}$$

For all o, $\tau_{oo} = 0$ (migration cost is zero if remaining at origin) and $R_{ot} = 1$ (origin earnings are in origin currency). We assume preferences q_{id} are distributed multivariate Fréchet with shape parameter θ, as in Bryan and Morten (2019).[7] This parameter determines the dispersion of preferences across locations. Let π_{dost} be the fraction of people of skill s from origin o choosing to work in d. Through the properties of the Fréchet distribution, this share can be written as:[8]

$$\pi_{dost} = \frac{(w_{dst} R_{dt}(1-\tau_{dot})\epsilon_{dot})^{\theta}}{\sum_k (w_{kst} R_{kt}(1-\tau_{kot})\epsilon_{kot})^{\theta}} \tag{B.2}$$

Intuitively, the share of individuals of skill s migrating from origin o to destination d is increasing in the destination wages in Philippine pesos, $w_{dst}R_{dt}$.

B.2.2 Migrant Income Shock and the Shift-Share Variable

Our model derives the shift-share variable that is our primary independent variable, making our model entirely consistent with our empirical framework.

We assume there are two skill groups in the population: high-skilled h and unskilled u ($s = \{h, u\}$).[9] At baseline ($t = 0$), the share of high-skilled and unskilled workers in province o are denoted, respectively, ℓ_{oh0} and ℓ_{ou0}, with $\ell_{ou0} = 1 - \ell_{oh0}$. Province-level global income per capita Y_{ot} depends on the distribution of worker locations and skill levels:

$$Y_{ot} = \sum_{s=h,u} \left[\ell_{ost} \sum_{d} (\pi_{dost} w_{dost} R_{dt}) \right] \tag{B.3}$$

Our shift-share variable isolates exogenous variation in only the migrant income portion of Y_{ot}, due to the 1997 exchange rate shocks. Let $\tilde{\Delta}$ refer to a short-run change. $\tilde{\Delta} R_d$ is the short-run change in destination d exchange rate.[10]

The short-run migrant income change due to exchange rate shocks $\tilde{\Delta} R_d$ in province o

[7] Here, θ is the elasticity of migration with respect to the destination wage. In the standard formulation: $F(q_1, ..., q_D) = exp\left\{ - \left[\sum_{d=1}^{D} q_d^{-\theta} \right] \right\}$. The Fréchet assumption, while widely used in the migration literature (e.g., Bryan and Morten (2019); Tombe and Zhu (2019)) relies on an IIA assumption. An alternative would be to separate the decision to emigrate from the location choice. In our setting where international migration is fairly common (7.5% of households had a migrant abroad), and recruitment agencies facilitate migration, we think the Fréchet assumption is a reasonable approximation.

[8] Full derivations are in Appendix B.3.

[9] We micro-found the education decisions in Appendix B.3.2.

[10] In practice, we use the short-run 1997-1998 change following the July 1997 crisis to construct the shift-share variable. To signify this captures a short-run change, we include no subscript t in terms involving $\tilde{\Delta}$. Focusing on a shift-share variable capturing a short-run change is desirable because the immediate post-Crisis exchange rate changes are more plausibly exogenous than subsequent, longer-run exchange rate changes that may be endogenous to post-Crisis economic policies in destinations. We discuss this further in Subsection 2.4.2.1.

depends on the share of workers in each destination for each skill level.[11]

$$\tilde{\Delta}Y_o = \sum_{s=h,u} \left[\ell_{os0} \sum_d \left(\pi_{dos0} w_{dos0} \tilde{\Delta} R_d \right) \right] \equiv Shiftshare_o \tag{B.4}$$

In the pre-shock period ($t = 0$), let total population in an origin be Pop_{o0}, and the number of workers by skill be L_{os0}. Also, let the number of workers going from o to destination d be L_{dos0}, so that $\ell_{os0} \equiv \frac{L_{os0}}{Pop_{o0}}$, and $\pi_{dos0} \equiv \frac{L_{dos0}}{L_{os0}}$. Let w_{dos0} be average pre-shock income in destination d for workers of skill s from origin o.

The "exposure weight" ω_{do0}, serves as the "share" in the shift-share. As in the main paper, we define this as province o's pre-shock aggregate migrant income from destination d (summed across skill groups), divided by province population to yield a *per capita* variable: $\omega_{do0} \equiv \frac{\sum_{s=h,u} L_{dos0} w_{dos0}}{Pop_{o0}}$. Now rewrite Equation B.4:

$$Shiftshare_o = \sum_{s=h,u} \sum_d \frac{L_{os0}}{Pop_{o0}} \frac{L_{dos0}}{L_{os0}} w_{dos0} \tilde{\Delta} R_d = \sum_d \left(\omega_{do0} \tilde{\Delta} R_d \right) \tag{B.5}$$

This is precisely the independent variable we use in our estimation.

B.2.3 Education Investments

Migrant income may drive educational investments at home, for instance, by easing liquidity constraints or changing the returns to schooling. In Appendix B.3.2 we micro-found changes to human capital under various scenarios, and derive how the change in the share of high-skilled workers h in origin o is:

$$\Delta \ell_{oht} = \frac{1}{\Psi} \Delta Y_o = \frac{1}{\Psi} \sum_{s=h,u} \left[\ell_{os0} \sum_d \left(\pi_{dos0} w_{dos0} \tilde{\Delta} R_d \right) \right] = \frac{1}{\Psi} \underbrace{\sum_d \omega_{do0}}_{MigInc_{o0}} \times \underbrace{\frac{\sum_d \omega_{do0} \tilde{\Delta} R_d}{\sum_d \omega_{do0}}}_{Rshock_o}, \tag{B.6}$$

where $\frac{1}{\Psi}$ captures the effect of the migrant income shock on skill share.[12] The regression result in column 3 of Table 2.4 is our quantitative estimate of this skill response. Below, we unpack the implications of these changing skill shares.

[11] The origin as a destination drops out as there are no exchange rate changes for the origin.

[12] In Appendix B.3.2 we derive changes to human capital with liquidity constraints, with no liquidity constraints, or with no borrowing. For certain models, Ψ captures the cost of education. We are agnostic about whether the education response is due to liquidity constraints or changing returns to education. Some combination of the two is possible, and has little bearing for our quantification.

B.2.4 Gravity Estimation of Migration Flows

Accounting for the impact of migrant income shocks first requires an estimate of impacts on migration itself. In our gravity equation, the Frechet parameter θ pins down the elasticity of migrant flows (from o to d) with respect to destination d wages. This determines subsequent location choices and migrant income. Taking logs of the gravity equation B.2 yields the estimating equation:

$$\log\ \pi_{dost} = \theta \log\ w_{dst} + \theta \log\ R_{dt} + \theta \log\left(1 - \tau_{dot}\right) - \log\ \left[\sum_k \left(w_{kst} R_{kt} (1 - \tau_{kot}) \epsilon_{kot}\right)^{\theta}\right] + \theta \epsilon_{dot} \tag{B.7}$$

To estimate θ, we leverage the exogenous exchange rate shocks. The coefficient on $\log R_{dt}$ identifies θ. We implement this at the origin-destination-skill level using a differenced regression.[13]

$$\Delta log\ \pi_{dos} = \gamma_{os} + \theta \Delta log\ R_d + \tilde{\epsilon}_{dos}$$

Here, the Δs are the change between before and after the shock; and so this differenced regression is equivalent to including destination fixed effects. We further include the origin-by-skill fixed effects and cluster our standard errors at the destination level. The results are in Appendix Table B.9. We estimate $\theta = 3.42$.

B.2.5 Change in Migrant Flows: Predictions and Decomposition

Migration flows from origin o to destination d depend on the probability of migrating by skill level, and share of workers who are of each skill level: $\pi_{doht}\ell_{oht} + \pi_{dout}\ell_{out}$. Changes in wages both abroad (say, via exchange rates), and at home (say, via more entrepreneurial investment), will determine migration flows. The change in aggregate outflows from an origin o has the following components:[14]

[13] As is common in such data, a large fraction of these units have no flows, and so we use a Poisson pseudo-maximum likelihood (PPML) estimator.

[14] The term $\chi_o \equiv \theta \sum_{s=h,u} \ell_{ost} \left[(1 - \pi_{oost}) \sum_{d \neq o} \left(\pi_{dost} \frac{\Delta R_{dt}}{R_{dt}}\right) - \pi_{oost} \left(\pi_{oost} \frac{\Delta w_{ost}}{w_{ost}}\right)\right]$ captures second-order equilibrium adjustments. We measure and include it in all accounting exercises. Intuitively, changes in wages at home or exchange rates in destinations indirectly affect the choice of specific destinations. For instance, if the US exchange rate changes favorably, it would lead to more outflows, and if the Malaysian exchange rate changes unfavorably, there will be less emigration. Since both sets of exchange rates change simultaneously, a portion of the lower Malaysian emigration is redirected to the increase in US emigration. Equation B.36 shows a version with these indirect effects. The derivation is in Appendix B.3.4.

$$\Delta\ Flows_{ot} = \underbrace{\Delta\ell_{oht}\sum_{d\neq o}(\pi_{doh0}-\pi_{dou0})}_{\text{Education channel in outflows}} + \underbrace{\theta\sum_{d\neq o}(\ell_{oh0}\pi_{doh0}+\ell_{ou0}\pi_{dou0})\frac{\Delta R_{dt}}{R_{d0}}}_{\text{Exchange rate channel in outflows}}$$
$$-\underbrace{\theta\left(\ell_{oh0}\pi_{ooh0}\frac{\Delta w_{oht}}{w_{oh0}}+\ell_{ou0}\pi_{oou0}\frac{\Delta w_{out}}{w_{ou0}}\right)}_{\text{Domestic income stemming outflows}} - \underbrace{\chi_o}_{\text{Indirect re-sorting}} \tag{B.8}$$

First, the skilled and unskilled have different migration probabilities. If the skilled are more likely to migrate, then an increase in the fraction skilled will raise migration. If, alternatively, most jobs abroad are unskilled, then migration probabilities may fall. The effect of education on flows is captured by the first term, which is a product of two components: the education response $\Delta\ell_{oht}$, and skill-differential in migration probabilities $\pi_{doht} - \pi_{dout}$. Second, as exchange rates change favorably, there will be a migration response to higher compensation. This depends on θ (the elasticity of migration with respect to destination wages), the shock size ΔR_{dt}, and migration probabilities $\ell_{oht}\pi_{doht} + \ell_{out}\pi_{dout}$. This second term is the "Exchange rate channel in outflows." Finally, the shock can change local earning levels, affecting Δw_{ost}. For instance, earnings from abroad may fund investments in firms and household enterprises at origin locations. Increases in domestic income stem the outflow of migrants, as captured by this last channel, which again depends on the location elasticity with respect to wages θ. These components are each increasing functions of the exchange rate shocks, and suggest (as we test empirically) that the shock may change migrant flows. For instance, the first term ("Education channel in outflows") can be seen from Equations B.6 and B.8 to be:

$$\Delta\ell_{oht}\sum_{d\neq o}(\pi_{doh0}-\pi_{dou0}) = \frac{1}{\Psi}\underbrace{\sum_{d\neq o}(\pi_{doh0}-\pi_{dou0})}_{\text{Skill bias in outmigration}}\left(\underbrace{\sum_{d}\omega_{do0}}_{MigInc_{o0}}\times\underbrace{\frac{\sum_d \omega_{do0}\tilde{\Delta}R_d}{\sum_d \omega_{do0}}}_{Rshock_o}\right) \tag{B.9}$$

We use this framework to quantify the importance of the education and exchange rate channels. To quantify the education channel, we obtain (a) the education response to the income shock $\Delta\ell_{oht}$ from column 3 of Table 2.4, and obtain (b) the skill-differential in migration probabilities $\pi_{doh0} - \pi_{dou0}$ from the raw data. Figure B.2a shows that for every province, the likelihood of becoming an overseas worker is higher when the worker has more education. Therefore, increases in education should increase the flow of migrants from all provinces.

The role played by the exchange rate and wage channels is jointly determined by simultaneous changes to exchange rates across potential migration destinations (ΔR_{dt}) and increases in domestic wages Δw_{ost}. We obtain the increases in domestic wages for different skill groups from columns 1 and 2 of Appendix Table B.10. Migration responses to these, in turn, depend on the Frechet parameter θ, estimated in section B.2.4. We combine these estimates with measures of the shares of skilled and unskilled at each province, and propensity to migrate abroad by skill group at baseline to calculate the second and third terms in Equation B.8.

Together, these channels predict outflows. We validate the structure of our model by comparing model predicted flows to the OLS prediction from column 4 of Appendix Table B.10 in Appendix Figure B.3a. The strong upward sloping relationship indicates that the model does a good job of predicting migration flows. A number of provinces with a high predicted flow lie above the 45-degree line, suggesting that there may be other changes in those provinces or non-linearities in the empirical relationship between flows and migrant income changes.

Finally, we quantify the role played by each channel. We calculate the share of the total regression-based predicted flows attributable to the education channel: $\frac{\Delta \ell_{oht} \sum_d (\pi_{doh0} - \pi_{dou0})}{\widehat{Flows}_{ot}^{OLS}}$. Appendix Figure B.3b plots the distribution of the contribution of the education channel across provinces. On average about 17.2% of the increase in migrant flows is attributable to the increased education response (Table B.11).[15] We do a similar exercise for the exchange rate channel. The exchange rate changes abroad will tend to drive migration abroad as most exchange rates changed favorably relative to the Philippines. At the same time, however, improvements in domestic income stem such outflows, canceling out a large component of the gains from migration. On net, changes in relative prices explain about 29.7% of the outflows. The remaining half is unexplained. We may not expect to explain the entire flows as we use baseline (1995) shares of migration flows.

B.2.6 Change in Migrant Income: Predictions and Decomposition

The change in migrant income per capita can be decomposed into: (1) the education channel, and (2) the persistent change in exchange rates, which raises migrant income and encourages

[15] Theoretically, the education channel contribution can be negative if the low-skilled have a higher migration probability.

flows to favorable destinations.

$$\underbrace{\Delta\ell_{oht}\left(\sum_{d\neq o} w_{doh0}\pi_{doh0}R_{d0} - \sum_{d\neq o} w_{dou0}\pi_{dou0}R_{d0}\right)}_{\text{Education channel in migrant income}} + \underbrace{\theta\left(\sum_{s=h,u}\left[\ell_{os0}\sum_{d}\left(\pi_{dos0}w_{dos0}\Delta R_{dt}\right)\right]\right) - \tilde{\chi}_{o2}}_{\text{Exchange rate channel in migrant income}} \tag{B.10}$$

Here, we know $\Delta\ell_{ost}$ is a function of the migrant income shock from Equation B.6. We define $\beta^{mig} = \left(\sum_{d\neq o} w_{doh0}\pi_{doh0}R_{d0} - \sum_{d\neq o} w_{dou0}\pi_{dou0}R_{d0}\right)$ as the migrant skill premium. The education channel contribution to the change in income is simply $\frac{\beta^{mig}}{\Psi}\tilde{\Delta}Y_o$. Similarly, the exchange rate channel is simply $\theta\tilde{\Delta}Y_o - \tilde{\chi}_{o2}$, and captures the increase in long run migrant income, not simply due to the fact that better exchange rates directly increase migrant income, but also because they induce a higher flows of migrants (both skilled and unskilled) to places with more positive exchange rate movements.[16] Additionally, as captured by what we call 'indirect resorting,' simultaneous changes in the exchange rate affect the location choices of migrants, which in turn affects how much they earn. The total change in migrant income per capita $\left(\frac{\beta^{mig}}{\Psi} + \theta\right)\tilde{\Delta}Y_o - \tilde{\chi}_{o2}$ is empirically shown in Table 2.3 col 5.

To quantify the importance of each component, we decompose the contributions of each channel. For the education channel, we first obtain $\Delta\ell_{ost}$ with the help of linear fit of the regression in column 3 of Table 2.4. The second component is the probability-weighted skill-premium abroad β^{mig}. We plot the skill premium $(w_{doh0} - w_{dou0})$ at the origin-destination pair in Figure B.2b.[17]

For the exchange rate channel, we use our estimate of θ. A higher migration elasticity θ means that migration flows, and thereby migrant income, are more responsive to exchange rate shocks. We measure the shares ℓ_{os0} and π_{dos0}, and wages w_{dos0} at baseline (1995), and use them as weights for exchange rate changes ΔR_{dt} as in the second term of Equation B.10.

Together, the predicted migrant income estimate due to the education channel and the exchange rate channel can be compared to the simple OLS prediction based on the regression from column 5 of Table 2.3. We plot the relationship between these predicted flows in Figure B.4a. As before, we see a strong upward sloping relationship in Figure B.4a which indicates that the model does a good job of predicting migrant income per capita. Predicted values are distributed around the forty-five degree line.

To quantify the role played by each channel, we measure the predicted education channel

[16] As before, the second-order indirect effects of changes in location choice are captured by $\tilde{\chi}_{o2} \equiv \theta\sum_{s=h,u}\sum_d\left[\ell_{ost}w_{dst}\pi_{dost}\left(\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right) + \pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right)\right]$.

[17] Returns are weighted by migration probabilities, as for many low-skilled occupations there are no migrant opportunities for certain destinations. As such, increases in skill raise earning prospects by raising employment prospects.

as a ratio of the predicted increase in migrant incomes (Appendix Figure B.4b). We do a similar exercise for the exchange rate channel in migrant income. On average, the education channel explains 24.4% of the increase in migrant income, and the exchange rate channel explains 75.5% (Table B.11).[18]

B.2.7 Change in Domestic Income: Prediction and Decomposition

Domestic income can rise for at least two reasons. First, an increase in education and skills allows workers to work in high-paying skilled jobs (the "Education channel"). Second, earnings from domestic work (conditional on skill) may also increase as a result of more local investment in enterprises and an increase in aggregate demand (the "Direct wage channel"). While simple to introduce, we do not explicitly model firm production to keep our framework simple and tractable. While the underlying mechanisms are not modeled, our framework captures the ultimate affect of the shock on domestic earnings. Specifically, investments in entrepreneurial capital and aggregate demand will raise domestic income for each skill group Δw_{ost}, and investments in human capital will raise the share high-skilled $\Delta \ell_{oht}$. Together, these increase domestic income per capita:

$$\Delta W_{ot} = \underbrace{\Delta \ell_{oht} \left(\underbrace{w_{oh0}\pi_{ooh0}}_{\text{skilled wage at home}} - \underbrace{w_{ou0}\pi_{oou0}}_{\text{unskilled wage at home}} \right)}_{\text{Education channel in domestic income}} + \underbrace{\sum_{s=h,u} \ell_{os0}\pi_{oos0} \left(\Delta w_{ost}\right) - \tilde{\chi}_{o1}}_{\text{Direct wage (and resorting) channel}} \quad \text{(B.11)}$$

Here, the domestic "direct wage channel" captures the direct effect of changes in local wages due to, say, expansion of household entrepreneurship (and the indirect effects of staying back/or emigrating given the relative changes in wages at home and abroad).[19] As we do not take a stance on the mechanisms underlying enterprises decisions, we allow Δw_{ost} to be a function of migrant income per capita. As we show in Section B.2.6, migrant income per capita is a function of the exchange rate shock: $\left(\frac{\beta^{mig}}{\Psi} + \theta\right) \tilde{\Delta} Y_o$. Let ζ be a local multiplier driven by changes to aggregate demand and entrepreneurial investments. In that

[18] It is not unreasonable for our model to explain a little more than the entirety of the changes, as we use baseline earnings in various destinations that may change for reasons unrelated to the shocks.

[19] The indirect resorting is $\tilde{\chi}_{o1} \equiv \sum_{s=h,u} \ell_{ost}\theta\pi_{oost}w_{ost}\left(\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right) + \pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right) - \sum_{s=h,u} \ell_{ost}\pi_{oost}\theta\Delta w_{ost}$

case, $\Delta w_{ost} \equiv \zeta \left(\frac{\beta^{mig}}{\Psi} + \theta\right) \tilde{\Delta} Y_o$. We empirically estimate the associated regression:

$$\Delta W_{ot} = \underbrace{\sum_{s=h,u} \ell_{os0} \pi_{oos0} \left(\zeta \left(\frac{\beta^{mig}}{\Psi} + \theta\right) \tilde{\Delta} Y_o\right)}_{\text{Direct wage channel}} + \underbrace{\frac{1}{\Psi} \tilde{\Delta} Y_o \left(w_{oh0}\pi_{ooh0} - w_{ou0}\pi_{oou0}\right)}_{\text{Education channel in domestic income}}$$
$$= \left(\zeta \left(\frac{\beta^{mig}}{\Psi} + \theta\right) + \frac{\beta^{dom}}{\Psi}\right) \underbrace{\sum_d \omega_{do0}}_{MigInc_{o0}} \times \underbrace{\frac{\sum_d \omega_{do0} \tilde{\Delta} R_d}{\sum_d \omega_{do0}}}_{Rshock_o}, \tag{B.12}$$

where $\beta^{dom} \equiv (w_{oh0}\pi_{ooh0} - w_{ou0}\pi_{oou0})$ are the domestic returns to education. We test for the change in domestic income per capita in Table 2.3 above.

We closely follow the methods described above for migrant income to again distinguish these channels. For instance, since the shock may directly change income at home, we use the baseline skill-premium when quantifying the education channel. Again, we aggregate predicted domestic income due to the education channel and the direct wage channel, and create a composite measure of predicted increases in domestic income per capita. We validate the model by comparing the model-predicted domestic income per capita with the simple OLS prediction based on the regression from column 4 of Table 2.3. We plot the relationship between these predicted flows in Appendix Figure B.5a. As before, we see a strong upward sloping relationship. The model slightly under-predicts domestic income per capita. Predicted values are distributed around the 45° line.

To quantify the role played by the direct wage channel, we estimate the impact of the migrant income shock on domestic income per worker by skill level in columns 1-2 of Table B.10. The increases in skill-specific domestic incomes are weighted by the baseline skill-shares in each province, and the probabilities that individuals do not emigrate conditional on their skill levels, as in Equation B.11.

Finally we measure the role played by the education channel in domestic income, as a ratio of the predicted increase in domestic income per capita. We plot this in Figure B.5b. We do a similar exercise for the direct wage channel. On average, the education channel explains 22.8% of the increase in domestic income, whereas the direct wage channel explains 60.8% (Table B.11). The remaining component is likely driven by other aggregate changes to the income distribution.

B.2.7.1 Explaining Impacts on Direct Domestic Income

In this section, we investigate the assumptions needed to explain the magnitude of the impact on domestic income per capita. As discussed in Subsection 2.6.3 of the main text, we need to explain how a 1 PhP migrant income shock leads to a 18.81 PhP increase in long-run domestic income, which is the coefficient estimate on the shift-share variable in the domestic income per capita regression of Table 2.3, Panel D column 4. 22.8% of the increase in domestic income can be attributed to the increase in education induced by the shock (as discussed in Section 2.6.1). This leaves the remaining 14.6 PhP increase to be explained. Here, we describe the framework in which we assess whether an effect of this size is reasonable.

We examine whether this remaining 14.6 PhP increase in domestic income per capita can be generated in a stylized framework in which a portion of the exogenous increase in migrant income is devoted to capital accumulation in productive enterprises, and in which a demand multiplier also operates. In every post-shock period t, an origin area enjoys the following increment to income per capita (we suppress origin o subscripts for simplicity):

$y_t = \alpha m_t + r_t S_{t-1}$, where m_t is exogenous migrant income per capita, α is the share of migrant income that is spent in the origin economy, S_t is the induced savings in the economy due to the shock, and r_t is the return to capital.

An exogenous portion s of the additional income is saved (and invested) each period, with shock-induced savings accumulating as: $S_t = S_{t-1} + sy_t$.

The shock-induced increase in domestic income per capita is then simply the shock-induced incremental per period income (y_t) multiplied by the Keynesian multiplier ($\frac{1}{s}$). We set the savings rate to 0.35, which implies a Keynesian multiplier of 2.86 (comparable to the 2.9 estimate in Breza and Kinnan (2021)). For migrant income m_t, given we are interested in the result of a 1 PhP shock, we set the initial shock $m_1 = 1$ and let the shock to evolve according to a function that asymptotically reaches our migrant income coefficient for 2015 ($m_\infty = 6.3$), and passes through our migrant income coefficient for 2009 ($m_{12} = 4.9$) from the event study (Figure 2.3).

We set the rate of return to initial rate $r_1 = 0.45$; this is high, but not as high as the estimate of de Mel et al. (2008). We then let r_t decline over time, according to a function that asymptotically reaches 0.05. This decline captures that the initial rate of return to capital may be quite high when liquidity constraints on investment are first loosened, but r_t declines over time as the most profitable investment opportunities are taken.[20]

Appendix Figures B.7a and B.7b trace out the shock-induced domestic income generated under these assumptions. The remaining 14.6 PhP increase in migrant income per capita is fully explainable, and is well within plausible assumptions. See the main text for discussion.

[20]The functional forms for the path of migrant income and rate of returns on savings are as follows: $m_t = \frac{6.32t^2-1.95t-0.37}{t^2+3t}$ and $r_t = \frac{0.05t^2+0.85t}{t^2+t}$. Time t is relative to 1997, where $t = 1$ is for 1998, and so on.

B.2.8 Change in Global Income: Predictions and Decomposition

Together, the longer-term change in the global income of individuals is:[21]

$$\left(\frac{\beta^{mig}+\beta^{dom}}{\Psi}+\theta+\zeta\left(\frac{\beta^{mig}}{\Psi}+\theta\right)\right)\tilde{\Delta}Y_o-\tilde{\chi}_o \tag{B.13}$$

There is intuition behind this relationship.[22] First, higher skill-premia (the β terms) imply that as individuals acquire schooling, incomes (both domestic and international) rise. Second, a higher migration elasticity θ means that migration flows, and thereby migrant incomes, are more responsive to favorable exchange rates. Finally, if incomes rise locally, then that would have a direct impact on income as well. Local incomes may rise through increases in aggregate demand or entrepreneurial investment, for instance.

In the long run, global income and household expenditure increase substantially, as we show in column 3 of Table 2.3. Overall changes in expenditure (column 4 of the same table) reflect changes in welfare. As we show, our theoretical predictions are consistent with our empirical predictions. This allows us to interpret our reduced form estimates, rationalize the magnitudes, and quantify the contribution of each channel discussed.[23]

Together, the changes in migrant income and domestic income allow us to decompose the changes in global income per capita. To test the validity of the model, we again predict the change the global income per capita using the regression estimated in column 3 of Table 2.3 for global income. Appendix Figure B.6a shows that our model again does a good job of predicting the change in global income. Since the domestic and migrant income channels both have an education component, we can again measure the total contribution of education investments to changes in global income. Figure B.6b plots the distribution of this contribution across provinces. Table B.11 shows that the education channel explains 23.2% of the overall increase in global income, while the changes in earnings potential (both at home and abroad) explain 64.2% of the overall increase in global income. Overall, the model explains 87.3% of the increase in global income.

[21] The derivation for global income is in Appendix B.3.5.

[22] The total indirect effect on global income due to location resorting is $\tilde{\chi}_o \equiv \theta\sum_{s=h,u}\sum_d\left[\ell_{ost}w_{dst}\pi_{dost}\left(\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right)+\pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right)\right]-\theta\sum_{s=h,u}\left[\ell_{ost}\pi_{oost}\Delta w_{ost}\right]$

[23] A short note on the model equilibrium. While simple to introduce, we do not explicitly model production to keep the analysis tractable and self-contained. Changes in production, whether at large firms or household enterprises, will affect domestic wages, changes to which are captured in our framework. Furthermore, this is not a spatial model of bilateral flows, where origins can be destinations and vice versa. With bounded migration costs, and a lack of agglomeration or congestion forces, we expect that labor and output markets clear in equilibrium (Allen et al., 2020).

B.3 Model Derivations

B.3.1 Deriving share of flows from o to d

Indirect utility of worker i is as defined in the text:

$$V_{idost} = w_{dst} R_{dt} (1 - \tau_{dost}) q_{id} \epsilon_{dot} \equiv \widetilde{w_{dost}} q_{id} \tag{B.14}$$

Workers will pick the destination p with the highest value of $w_{idost} = \widetilde{w_{dost}} q_{id}$. The probability that they pick destination 1 is given by:

$$\begin{aligned}
\pi_{1ost} &= Pr\left[\widetilde{w_{1ost}} q_1 > \widetilde{w_{d'ost}} q_{d'}\right] \quad \forall d' \neq 1 \\
&= Pr\left[q_{d'} < \frac{\widetilde{w_{1ost}} q_1}{\widetilde{w_{d'ost}}}\right] \quad \forall d' \neq 1 \\
&= \int \frac{dF}{dq_1} \left(q_1, \alpha_2 q_1, ..., \alpha_D q_1\right) dq_1
\end{aligned} \tag{B.15}$$

where we define $\alpha_d \equiv \frac{\widetilde{w_{1ost}}}{\widetilde{w_{d'ost}}}$. We assume that the abilities are distributed with the following Frechet distribution:

$$F(q_1, ..., q_D) = exp\left\{ - \left[\sum_{d=1}^{D} q_d^{-\theta} \right] \right\} \tag{B.16}$$

So the derivative of the CDF is given by:

$$\frac{dF}{dq} = \theta q^{-\theta-1} exp\left\{ - \left[\sum_{d=1}^{D} q_d^{-\theta} \right] \right\} \tag{B.17}$$

This derivative evaluated at $(q_1, \alpha_2 q_1, ..., \alpha_D q_1)$, allows us to determine the probability of choosing destination 1:

$$
\begin{aligned}
\pi_{1ost} &= \int \theta q^{-\theta-1} exp\left\{-\left[\sum_{d=1}^{D}(\alpha_d q)^{-\theta}\right]\right\} dq \\
&= \frac{1}{\sum_{d=1}^{D}\alpha_d^{-\theta}} \int \left(\sum_{d=1}^{D}\alpha_d^{-\theta}\right) q^{-\theta-1} exp\left\{-\left[q^{-\theta}\left(\sum_{d=1}^{D}\alpha_d^{-\theta}\right)\right]\right\} dq \\
&= \frac{1}{\sum_{d=1}^{D}\alpha_d^{-\theta}} \int dF(q) \\
&= \frac{1}{\sum_{d=1}^{D}\alpha_d^{-\theta}} .1 \\
&= \frac{\widetilde{w_{1ost}}^{\theta}}{\sum_{d=1}^{D}\widetilde{w_{dost}}^{\theta}}
\end{aligned} \tag{B.18}
$$

The third line comes from the properties of the Frechet distribution, where we know that the term in the integral of the second line is simply the PDF with a shape parameter θ, and a scale parameter $\sum_{d=1}^{D}\alpha_d^{-\theta}$. Expanding on the definitions for $\widetilde{w_{dost}}$, and including the subscripts, we get equation B.2:

$$
\pi_{dost} = \frac{(w_{dst}R_{dt}(1-\tau_{dot})\epsilon_{dot})^{\theta}}{\sum_k (w_{kst}R_{kt}(1-\tau_{kot})\epsilon_{kot})^{\theta}} \tag{B.19}
$$

B.3.2 Micro-founding the Education Responses

Baseline Framework: Households choose schooling levels S when young, and how much to borrow b_{io}. They maximize two period utility: $u(c_1)+u(c_2)$. Period 1 consumption depends on wealth Y (including migrant income), the price of schooling p, and borrowing. Period 2 consumption depends on income and period 1 debt with interest I:

$$
c_{1io} = Y_{io} - p_o S_{io} + b_{io} \quad \text{and} \quad c_{2io} = V_{idost} - I_o b_{io} \ , \tag{B.20}
$$

where w_{idost} is the wage after the location choice.

We may expect that changes in migrant income help drive investments in human capital at home, for instance, by easing liquidity constraints for households or changing the returns to schooling. For instance, under certain assumptions on $u(.)$ and w say, $w_{do}(S)$ linear in S, and log-utility $u(c)$ and for credit constrained households $\bar{b}=0$, average province-level schooling responds to shocks to migrant income: $\Delta S_{ot} = \frac{1}{2p}\Delta Y_o$. In this case, for $\Psi \equiv (ed_1 - ed_0)2p$, the change in the share of high-skilled workers h in origin o is:

$$\Delta \ell_{oht} = \frac{1}{\Psi}\Delta Y_o = \frac{1}{\Psi}\sum_{s=h,u}\left[\ell_{os0}\sum_d \left(\pi_{dos0} w_{dos0}\Delta R_{d0}\right)\right] = \frac{1}{\Psi}\underbrace{\sum_d \omega_{do0}}_{MigInc_{o0}} \times \underbrace{\frac{\sum_d \omega_{do0}\tilde{\Delta} R_d}{\sum_d \omega_{do0}}}_{Rshock_o} \quad \text{(B.6)}$$

Non Credit Constrained Households and Changes in Returns: Non constrained households may also respond to exchange rate shocks. Exchange rate shocks may not change the returns to education as they change both the educated and non-educated wage. For those who are not constrained, we derive that for a cost of education $= p_1 S + p_2 S^2$, the optimal amount of schooling does not depend on Y, but only on the returns to education:

$$S_i^u = \frac{w'(s)_d(1-\tau_{dot})R_{dt}q_{id} - p_1}{2p_2} \quad \text{(B.21)}$$

where S_i^u are the years of schooling for unconstrained households. The average education levels of non-constrained households from origin o to destination d are:

$$S_{do}^u = \frac{w'(s)_d(1-\tau_{dot})R_{dt}\pi_{dot}^{\frac{-1}{\theta}}\Gamma - p_1}{2p_2} \quad \text{(B.22)}$$

And the average change in education for unconstrained households from origin o is:

$$S_o^u = \sum_d S_{do}\pi_{dot} = \sum_d \frac{w'(s)_d(1-\tau_{dot})R_{dt}\pi_{dot}^{\frac{-1}{\theta}+1}\Gamma - p_1}{2p_2} \quad \text{(B.23)}$$

Since $\Delta \pi_{dot}^{\frac{-1}{\theta}} = -\pi_{dot}^{\frac{-1}{\theta}}\frac{\Delta R_{dt}}{R_{dt}}$, we know that:

$$\Delta S_o^u = \sum_d \frac{w'(s)_d(1-\tau_{dot})\theta\pi_{dot}\Gamma}{2p_2}\frac{\Delta R_{dt}}{R_{dt}} \quad \text{(B.24)}$$

If δ fraction of the population is credit constrained, then the education response will also depend on δ. Notice that for unconstrained households to respond, students must also expect the exchange rate shocks to be long lasting.

Constraints on borrowing from future: For borrowing constrained households, the amount of schooling will depend on the income in the first period (and thereby any shocks to the income from abroad). Consider the two period consumption problem in Equation

B.20, and the lifetime utility $u(c_1) + u(c_2)$. If $b = \bar{b}$ is binding, then schooling is the only choice. From the first order conditions with respect to schooling, we know that:

$$pu'(c_1) = w'(S)u'(c_2) \tag{B.25}$$

For continuous, increase and concave utility and earnings functions, using the implicit function theorem, we can show education is an increasing function of income $\frac{\Delta S}{\Delta Y} > 0$.[24] We can also derive meaningful closed form solutions under other assumptions, such as for a linear earnings function: $w(S) = w'(S)S$, and Cobb-Douglas utility, say $u(c) = \alpha logc$, we can show that for $\bar{b} = 0$ (completely constrained households), the first order condition is simply: $\frac{p\alpha}{Y-pS} = \frac{\alpha}{w(S)}w'(S)$. We can derive a simple closed form relationship: $S_o = \frac{1}{2p}Y_o$.
For partially binding credit constraints, we can show $\Delta S = \frac{-I\bar{b}}{4p\gamma_d(1-\tau_{do})q_{id}R_{dt}}\frac{\Delta R_{dt}}{R_{dt}}$, where I is the rate of interest on borrowing

We are agnostic about whether the education response is due to liquidity constraints or changing returns to education. Some combination of the two is possible. Additionally, if period 2 consumption is subjectively discounted, say at rate β, then both the education and skill-share response will be scaled by $\frac{\beta}{1+\beta}$.

B.3.3 Deriving the changes in π_{dost}

Flows from origin o to destination d are given by Equation B.19. We define V_{ost} as the denominator of Equation B.19. That is, $V_{ost} \equiv \sum_k (w_{kst}R_{dt}(1-\tau_{kot})\epsilon_{kot})^\theta$. This comes to represent the option value of working in the various possible destinations. Similarly, let us define the numerator of Equation B.19 to be $V_{dost} = (w_{dst}R_{dt}(1-\tau_{dot})\epsilon_{dot})^\theta$.

$$\pi_{dost} = \frac{(w_{dst}R_{dt}(1-\tau_{dot})\epsilon_{dot})^\theta}{\sum_k (w_{kst}R_{kt}(1-\tau_{kot})\epsilon_{kot})^\theta} \equiv \frac{V_{dost}}{V_{ost}} \tag{B.19}$$

We can take the total derivative of these flows with respect to changes (derivative) in the exchange rate for one specific destination ΔR_{dt}:[25]

$$\Delta\pi_{dost} = \underbrace{\frac{((1-\tau_{dot})\epsilon_{dot})^\theta}{V_{ost}}\left(w_{dst}^\theta \theta R_{dt}^{\theta-1}\Delta R_{dt} + R_{dt}^\theta \theta w_{dst}^{\theta-1}\Delta w_{dst}\right)}_{\text{from the numerator of Equation B.19}} - \underbrace{\frac{V_{dost}}{V_{ost}^2}\Delta V_{ost}}_{\text{from the denominator of Equation B.19}} \tag{B.26}$$

[24]To be specific: $\frac{\Delta S}{\Delta Y} = p + \frac{u''(c_2)}{u''(c_1)}\frac{w'(S)}{p} + \frac{u'(c_2)}{u''(c_1)}\frac{w''(S)}{p}$. Since $u'(c) > 0$, $u''(c) < 0$, $w'(S) > 0$, $w''(S) < 0$, we know $\frac{\Delta S}{\Delta Y} > 0$.

[25]Here, and elsewhere, we use Δ to denote a derivative, as d is already used for destinations.

The above equation is derived using the quotient rule. The first part takes changes in the numerator, where only R_{dt} and w_{dst} change. This captures the effect of the exchange rate shocks to destination d specifically. Yet, simultaneously every exchange rate and every origin's wage changes as a result of the shock. So how does the π_{dost} change when there are multiple indirect changes as well? The second part takes the total derivative of the denominator. Now, since $\pi_{dost} \equiv \frac{V_{dost}}{V_{ost}}$, we can simplify this further:

$$\Delta\pi_{dost} = \underbrace{\theta\pi_{dost}\left(\frac{\Delta R_{dt}}{R_{dt}} + \underbrace{\frac{\Delta w_{dst}}{w_{dst}}}_{=0\ if\ o\neq d}\right)}_{\text{from the numerator of Equation B.19}} - \underbrace{\frac{\pi_{dost}}{V_{ost}}\Delta V_{dost}}_{\text{from denominator of Equation B.19}} \tag{B.27}$$

For all $d \neq o$ the shocks do not change destination wages (i.e. Filipino migrants are small enough a group in destinations to affect their equilibrium wages). As such, for such destinations, we know that there is a direct effect, and an indirect effect to go to specific destination d:

$$\Delta\pi_{dost} = \theta\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}} - \frac{\pi_{dost}}{V_{ost}}\left[\sum_{d\neq o}\left(V_{dost}\theta\frac{\Delta R_{dt}}{R_{dt}}\right) + \left(V_{oost}\theta\frac{\Delta w_{ost}}{w_{ost}}\right)\right] \tag{B.28}$$

This can be rewritten as:

$$\Delta\pi_{dost} = \theta\pi_{dost}\left[\underbrace{\frac{\Delta R_{dt}}{R_{dt}}}_{\text{Direct effect}} - \left(\underbrace{\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right)}_{\text{Indirect resorting}} + \underbrace{\pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}}_{\text{Domestic earnings stemming flows}}\right)\right] \tag{B.29}$$

Change in flows depends on shock on own destination, but also how flows would change to other destinations, and how increases to domestic income would stem such flows. This captures how flows to other destinations change, indirectly affect flows to the current destination.

We can sum up across destinations, and rewrite this equation

$$\sum_{d\neq o}\Delta\pi_{dost} = \theta\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\left[1 - \sum_{d\neq o}\pi_{dost}\right]\right) - \left(\theta\pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\left[\sum_{d\neq o}\pi_{dost}\right]\right) \tag{B.30}$$

$$\sum_{d\neq o} \Delta\pi_{dost} = \underbrace{\pi_{oost}\left[\theta \sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right)\right]}_{\text{Exchange rates driving outflows*}} - \underbrace{[1-\pi_{oost}]\left(\theta\pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right)}_{\text{Domestic earnings stemming outflows*}} \tag{B.31}$$

Alternatively, we could separate out the indirect sorting effects:

$$\begin{aligned}\sum_{d\neq o} \Delta\pi_{dost} = \quad & \underbrace{\theta \sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right)}_{\text{Exchange rates driving outflows}} \quad - \quad \underbrace{\theta\left(\pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right)}_{\text{Domestic earnings stemming outflows}} \\ & - \underbrace{\theta\left[\sum_{d\neq o}\pi_{dost}\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right) - \left[1-\sum_{d\neq o}\pi_{dost}\right]\pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right]}_{\text{Indirect resorting}}\end{aligned} \tag{B.32}$$

B.3.4 Deriving the changes in total flows

The above derivation is for a specific skill level s. Yet, skill levels may change as a result of the shock, and different skill groups have different propensities to migration. We know that flows from a specific origin to a specific destination can be characterized by:

$$\pi_{doht}\ell_{oht} + \pi_{dout}\ell_{out} \tag{B.33}$$

Suppose, only R_{dt} changed for one d, and there were no changes to domestic wages, then the direct effect would come from the first part of Equation B.29:

$$\Delta\ Flows_{dot} = \underbrace{\Delta\ell_{oht}\left(\pi_{doht}-\pi_{dout}\right)}_{\text{Education channel in flows}} + \underbrace{\theta\left(\ell_{oht}\pi_{doht}+\ell_{out}\pi_{dout}\right)\frac{\Delta R_{dt}}{R_{dt}}}_{\text{Exchange rate channel in direct flows}} \tag{B.34}$$

The second part above (exchange rate channel in direct flows) comes straight from the first part (direct effect) of Equation B.29 replaced into Equation B.33.

Equation B.31 allows us to derive $\Delta\ Flows_{ot} \equiv \sum_{d\neq o} \Delta\ Flows_{dot}$:

$$\Delta\ Flows_{ot} = \underbrace{\Delta \ell_{oht} \sum_{d\neq o} \left(\pi_{doht} - \pi_{dout}\right) +}_{\text{Education channel in outflows}} \quad \underbrace{\theta \sum_{d\neq o} \left(\ell_{oht}\pi_{ooht}\pi_{doht} + \ell_{out}\pi_{oout}\pi_{dout}\right) \frac{\Delta R_{dt}}{R_{dt}}}_{\text{Exchange rate channel in outflows (from Equation B.31 part 1)}}$$
$$\underbrace{-\,\theta \left(\ell_{oht}\left[1-\pi_{ooht}\right]\pi_{ooht}\frac{\Delta w_{oht}}{w_{oht}} + \ell_{out}\left[1-\pi_{oout}\right]\pi_{oout}\frac{\Delta w_{out}}{w_{out}}\right)}_{\text{Domestic earnings stemming outflows (from Equation B.31 part 2)}} \tag{B.35}$$

We can split up the exchange rate channel by skill group:

$$\Delta\ Flows_{ot} = \underbrace{\Delta \ell_{oht} \sum_{d\neq o} \left(\pi_{doht} - \pi_{dout}\right)}_{\text{Education channel in outflows}} \tag{B.36}$$
$$+\,\theta \left[\underbrace{\ell_{oht}\pi_{ooht} \sum_{d\neq o} \left(\pi_{doht}\frac{\Delta R_{dt}}{R_{dt}}\right)}_{\text{Exchange rate driving skilled outflows*}} + \underbrace{\ell_{out}\pi_{oout} \sum_{d\neq o} \left(\pi_{dout}\frac{\Delta R_{dt}}{R_{dt}}\right)}_{\text{Exchange rate driving unskilled outflows*}}\right]$$
$$-\,\theta \left[\underbrace{\ell_{oht}\left[1-\pi_{ooht}\right]\pi_{ooht}\frac{\Delta w_{oht}}{w_{oht}}}_{\text{Domestic earnings stemming skilled outflows*}} + \underbrace{\ell_{out}\left[1-\pi_{oout}\right]\pi_{oout}\frac{\Delta w_{out}}{w_{out}}}_{\text{Domestic earnings stemming unskilled outflows*}}\right]$$

Here, the channels above include the indirect re-sorting to the alternative destinations. Alternatively, we can keep the indirect re-sorting separate and use Equation B.32:

$$\Delta\ Flows_{ot} = \underbrace{\Delta\ell_{oht}\sum_{d\neq o}(\pi_{doht}-\pi_{dout})}_{\text{Education channel in outflows}} - \underbrace{\chi_o}_{\text{Indirect re-sorting}} \tag{B.8}$$

$$+\,\theta\left[\underbrace{\ell_{oht}\sum_{d\neq o}\left(\pi_{oht}\frac{\Delta R_{dt}}{R_{dt}}\right)+\ell_{out}\sum_{d\neq o}\left(\pi_{dout}\frac{\Delta R_{dt}}{R_{dt}}\right)}_{\text{Exchange rate driving outflows by skill group}}\right]$$

$$-\,\theta\left[\underbrace{\ell_{oht}\pi_{ooht}\frac{\Delta w_{oht}}{w_{oht}}+\ell_{out}\pi_{oout}\frac{\Delta w_{out}}{w_{out}}}_{\text{Domestic earnings stemming outflows by skill group}}\right]$$

where $\chi_o \equiv \theta\sum_{s=h,u}\ell_{ost}\left[(1-\pi_{oost})\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right)-\pi_{oost}\left(\pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right)\right]$

B.3.5 Contributions to changes in global income

The changes to income consist of two main components. First, let us look at domestic income (for those who do not migrate):

$$\sum_{s=h,u}\ell_{ost}\pi_{oost}w_{ost} \tag{B.37}$$

The direct effect on the domestic income would exist if wages increased $\Delta w_{ost}\neq 0$. The first is just the direct "wage channel" – higher wage rates imply higher domestic income. The second is driven by the fact that measured income rises only because education levels rise, and skilled workers are paid more.

$$\Delta W_{ot} = \underbrace{\Delta\ell_{oht}\left(\underbrace{w_{oh0}\pi_{ooh0}}_{\text{skilled wage at home}} - \underbrace{w_{ou0}\pi_{oou0}}_{\text{unskilled wage at home}}\right)}_{\text{Education channel in domestic income}} + \underbrace{\sum_{s=h,u}\ell_{os0}\pi_{oos0}\left(\Delta w_{ost}\right)-\tilde{\chi}_{o1}}_{\text{Direct wage (and resorting) channel}} \tag{B.11}$$

Overall income generated by the individuals that originate from these regions changes by more than simply the direct wage and education channels. This is because the location choices of individuals change as well, in response to lucrative exchange rates, and domestic wage increases. If domestic wages increase, then more people may remain behind locally, and earn at home: $\Delta\pi_{oost}$. We can return to Equation B.27, and set $d=o$, and $\Delta R_{ot}=0$.

But this time, $\Delta w_{ost} \neq 0$. So the analogue of Equation B.29 is given by:

$$\Delta\pi_{oost} = \theta\pi_{oost}\left[\underbrace{\frac{\Delta w_{ost}}{w_{ost}}}_{\text{Remainers}} - \underbrace{\left(\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right) + \pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right)}_{\text{Indirect resorting}}\right] \tag{B.38}$$

There is also the indirect effect once again. If wages do not increase at home, more workers may leave if exchange rates abroad become more favorable, reducing domestic income.

How does $\Delta\pi_{oost}$ contribute to domestic earning increases? We can replace the result for $\Delta\pi_{oost}$ above into Equation B.37, and derive the indirect resorting $\tilde{\chi}_{o1} \equiv \sum_{s=h,u}\ell_{ost}\theta\pi_{oost}w_{ost}\left(\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right) + \pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right) - \sum_{s=h,u}\ell_{ost}\pi_{oost}\theta\Delta w_{ost}$.

While this captures the domestic income gains, migrant income may change as well. Migrant income is given by:

$$\sum_{s=h,u}\ell_{ost}\sum_{d}\pi_{dost}w_{dost}R_{dt} \tag{B.39}$$

Again, changes to ℓ_{ost} (upskilling) will contribute to the education channel, as always:

$$\Delta\ell_{oht}\left(\underbrace{\sum_{d\neq o}w_{doht}\pi_{doht}R_{dt}}_{\text{skilled wage abroad}} - \underbrace{\sum_{d\neq o}w_{dout}\pi_{dout}R_{dt}}_{\text{unskilled wage abroad}}\right) \tag{B.40}$$

Now to get at how changes to exchange rates directly (and changes to local wages indirectly) affect flows, and thereby incomes, we need to go back to Equation B.29, which described how flows changed. To be specific, the effects on income due to more favorable exchange rates are driven by higher persistent income, and more flows abroad to avail of these favorable exchange rates. To a specific destination d, this is again given by:

$$\Delta\pi_{dost} = \theta\pi_{dost}\left[\underbrace{\frac{\Delta R_{dt}}{R_{dt}}}_{\text{Direct effect}} - \underbrace{\left(\sum_{d\neq o}\left(\pi_{dost}\frac{\Delta R_{dt}}{R_{dt}}\right) + \pi_{oost}\frac{\Delta w_{ost}}{w_{ost}}\right)}_{\text{Indirect resorting}}\right] \tag{B.29}$$

Again, the indirect resorting channel depends on the relative changes to exchange rates in other destinations. From Equation B.39, we can see that the changes to income are driven by (1) $\Delta\ell_{ost}$ (shown in Equation B.40), (2) $\Delta\pi_{dost}$ (shown in Equation B.29), and (3) just direct changes to ΔR_{dt} (say, in the short run). Since Equation B.40 already documents how changes to skill affect income, let us concentrate on (2) and (3) here:

$$\sum_{s=h,u} \ell_{ost} \sum_{d} \Delta\pi_{dost} w_{dost} R_{dt} + \sum_{s=h,u} \ell_{ost} \sum_{d} \pi_{dost} w_{dost} \Delta R_{dt} \tag{B.41}$$

Replacing the result from Equation B.29 in the first part of the equation above, we know:

$$\sum_{s=h,u} \ell_{ost} \sum_{d} \theta\pi_{dost} \frac{\Delta R_{dt}}{\cancel{R_{dt}}} w_{dost} \cancel{R_{dt}} - \tilde{\chi}_{o2} \quad + \quad \sum_{s=h,u} \ell_{ost} \sum_{d} \pi_{dost} w_{dost} \Delta R_{dt} \tag{B.42}$$

where $\tilde{\chi}_{o2} \equiv \theta \sum_{s=h,u} \sum_{d} \left[\ell_{ost} w_{dost} \pi_{dost} \left(\sum_{d \neq o} \left(\pi_{dost} \frac{\Delta R_{dt}}{R_{dt}} \right) + \pi_{oost} \frac{\Delta w_{ost}}{w_{ost}} \right) \right]$ is the indirect re-sorting (from Equation B.29). Rewriting this in terms of the initial shock ΔY_o:

$$\theta \underbrace{\sum_{s=h,u} \ell_{ost} \sum_{d} \pi_{dost} w_{dost} \Delta R_{dt}}_{\Delta Y_o = \text{Migrant Earnings Shock}} + \underbrace{\sum_{s=h,u} \ell_{ost} \sum_{d} \pi_{dost} w_{dost} \Delta R_{dt}}_{\text{Short run } \Delta c_{1o} = \Delta Y_o} - \tilde{\chi}_{o2} \tag{B.43}$$

So together the contribution of wages and exchange rate changes (not skill-upgrading) to longer-run changes in global income generated (and consumption Δc_{2o}) by individuals from these regions (whether they are located at home or abroad) is given by:

$$\underbrace{\sum_{s=h,u} \left[\ell_{ost} \pi_{oost} \left(\underbrace{\Delta w_{ost}}_{\text{Direct wage channel}} + \underbrace{\theta \Delta w_{ost}}_{\text{Remainers channel}} \right) \right]}_{\text{Domestic earnings due to firm-side responses}} - \tilde{\chi}_{o2} + \underbrace{\theta \underbrace{\left(\sum_{s=h,u} \ell_{ost} \sum_{d} \pi_{dost} w_{dost} \Delta R_{dt} \right)}_{\Delta Y_o = \text{Migrant Earnings Shock}}}_{\text{Earnings from Abroad: Exchange Rate Channel}} \tag{B.44}$$

B.4 Additional Tables and Figures

Figure B.1: Persistence of Exchange Rate Shock and Province-Destination Migrant Income

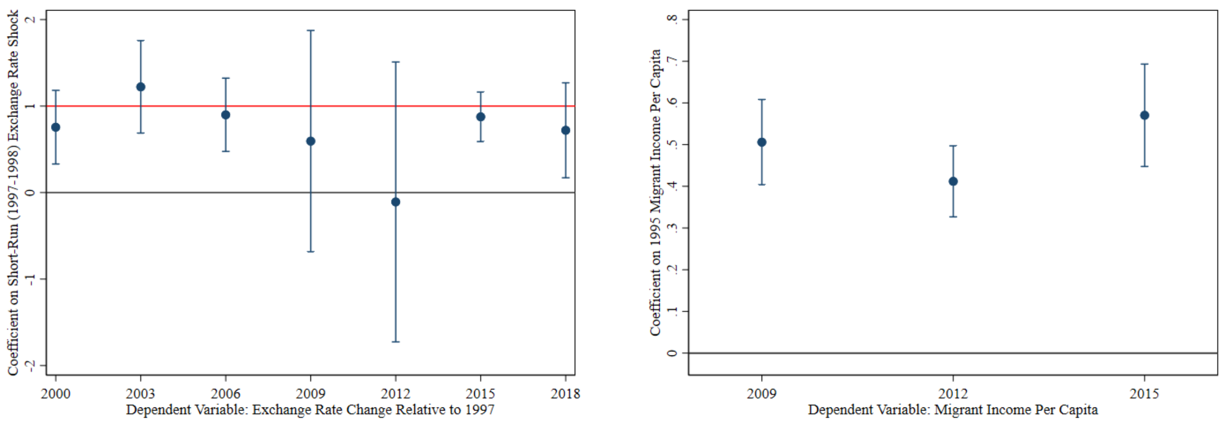

(a) $\tilde{\Delta} R_d$ and Future Exchange Rate Changes

(b) Province-Destination Migrant Income

Notes: (a) Coefficient estimates from regressing destination exchange rate changes relative to 1997 for 2000-2018 triennially on $\tilde{\Delta} R_d$, weighted by 1995 migrant income shares (N = 104). (b) Figure examines persistence from before to after the 1997 Asian Financial Crisis of ω_{dot} (migrant income per capita of province o from destination d). Figure displays coefficient estimates from regressing ω_{dot} for 2009, 2012, and 2015 (respectively) on ω_{do0} (1995 migrant income per capita, or the "exposure weight" used in the shift-share variable.) N = 74 $\times$ 104 = 7696, SEs clustered at province level.

Figure B.2: Skill Level, Migration Probabilities, and Migrant Wages

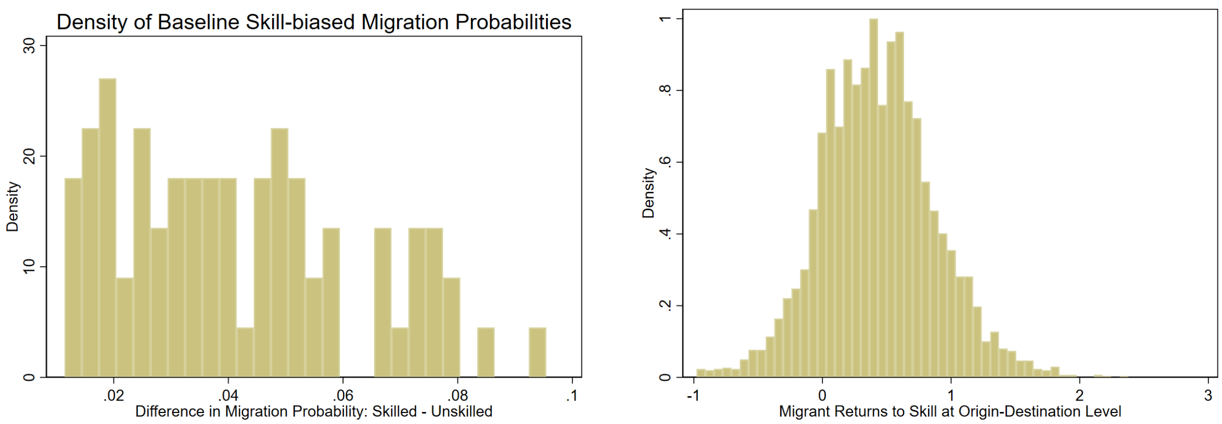

(a) Skilled-Unskilled Migration Probabilities (b) Wage skill-premium among migrants

Notes: (a) Figure plots a binned histogram of the difference in migration probabilities by skill, across provinces in 1990. We calculate the share of the skilled population that in the age-group 25-64 that is an overseas worker in destination d to be π_{dos}. We similarly do this for unskilled workers in π_{dou}. We then aggregate the difference across destinations, and plot $\sum_k (\pi_{kos} - \pi_{kou})$. (b) Figure plots the distribution of $w_{dost} - w_{dout}$ at the origin-destination pair level.

Figure B.3: Model Validation & Contribution of Education Channel in Migrant Flows

(a) Validation: Migrant flows (b) Contribution of Education Channel

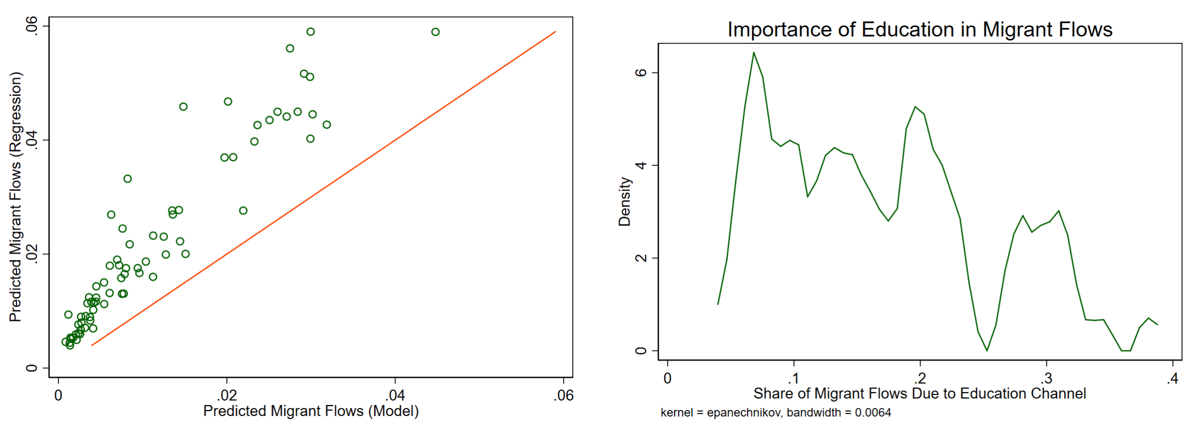

Notes: Figure B.3a plots the predicted flows of migrants vs the predicted flows as determined by the components of Equation B.8. The red line has an angle of 45 degrees. Each point represents a province. Figure B.3b plots the province-level distribution of the contribution of the education channel in predicting migrant flows: $\frac{\Delta \ell_{ost} \sum_k (\pi_{kos0} - \pi_{kou0})}{\widehat{Flows_{ot}^{OLS}}}$

Figure B.4: Model Validation & Contribution of Education in Migrant Income

(a) Validation: Migrant Income per capita

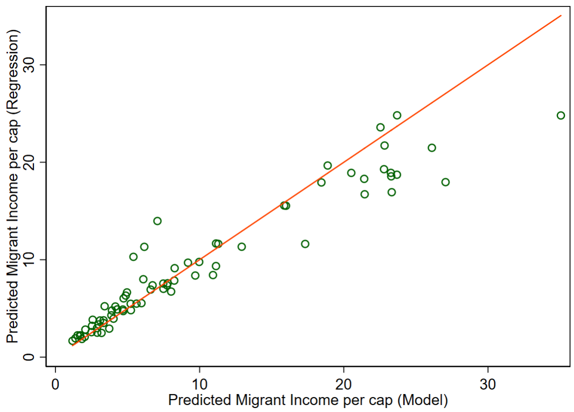

(b) Contribution of Education Channel

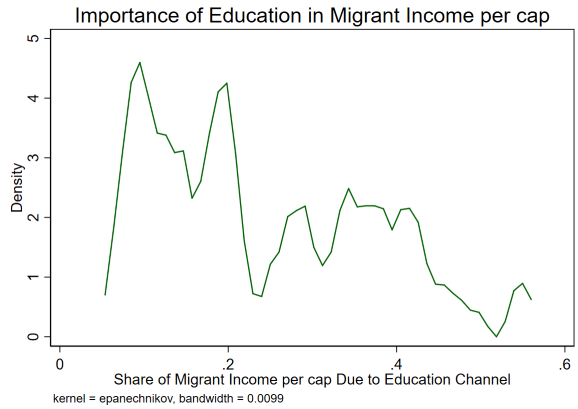

Notes: Figure B.4a plots the predicted migrant income per capita from the regressions (vertical axis) vs the predicted migrant income as determined by the education and exchange rate components. The red line has an angle of 45 degrees. Each point represents a province. Figure B.4b plots the province-level distribution of the contribution of the education channel in predicting migrant income per capita.

Figure B.5: Model Validation & Contribution of Education in Domestic Income

(a) Validation: Domestic Income per capita

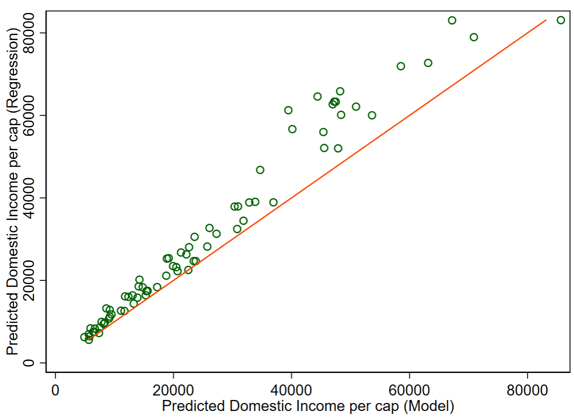

(b) Contribution of Education Channel

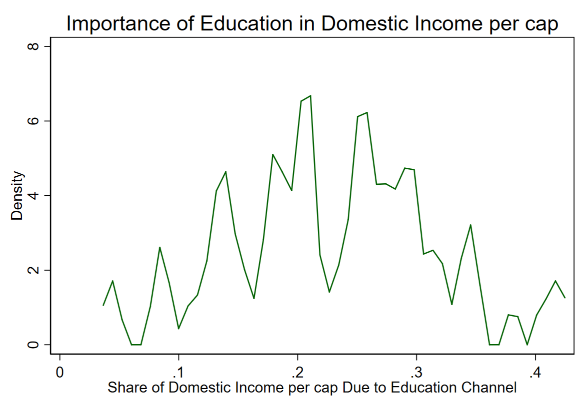

Notes: Figure B.5a plots the predicted domestic income per capita from the regressions vs the predicted domestic income per capita as determined by the education and exchange rate components. The red line has an angle of 45 degrees. Each point represents a province. Figure B.5b plots the province-level distribution of the contribution of the education channel in predicting domestic income per capita.

Figure B.6: Model Validation & Contribution of Education to Global Income

(a) Validation: Global Income per capita

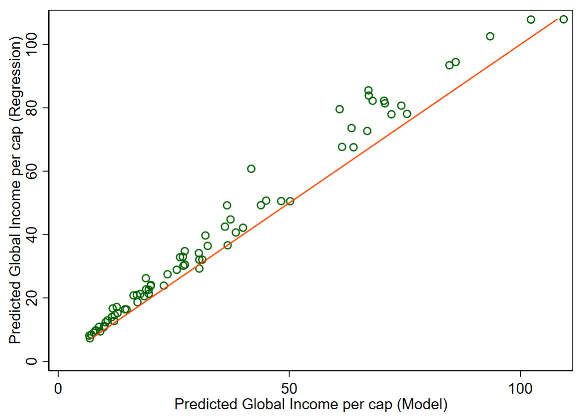

(b) Contribution of Education Channel

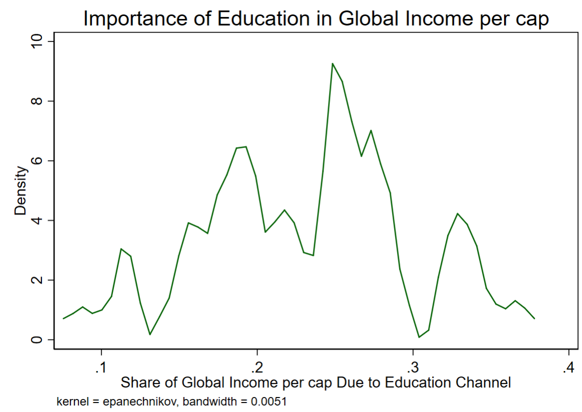

Notes: Figure B.6a plots the predicted global income per capita (domestic plus migrant income) from the regressions vs the predicted global income per capita as determined by the education and exchange rate components. The red line has an angle of 45 degrees. Each point represents a province. Figure B.6b plots the province-level distribution of the contribution of the education channel in predicting global income per capita.

Figure B.7: Explaining Effect on Domestic Income: Sensitivity to Key Assumptions

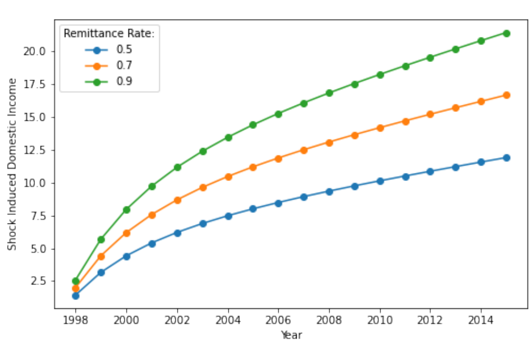

(a) Domestic Income Effects by Share of Migrant Income Spent at Origin (α)

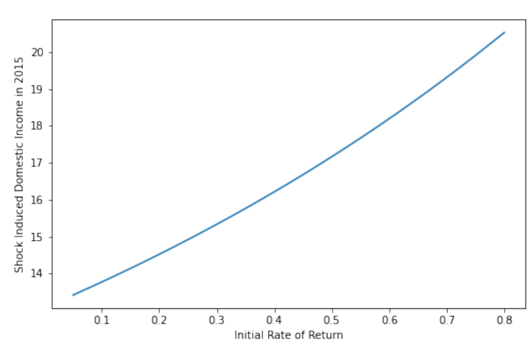

(b) Impact on Domestic Income by 2015, by Initial Rate of Return to Capital

Figure B.8: Event Studies for Other Outcomes

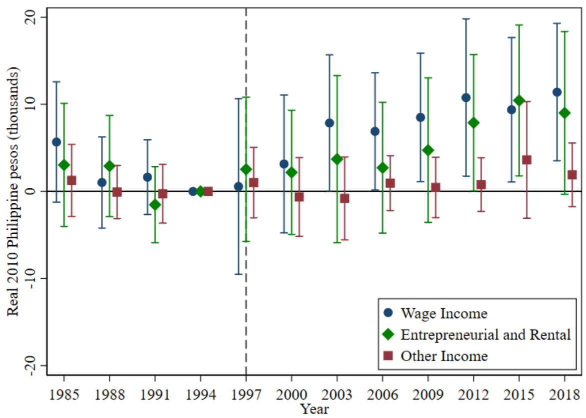

(a) Domestic Income Subcomponents

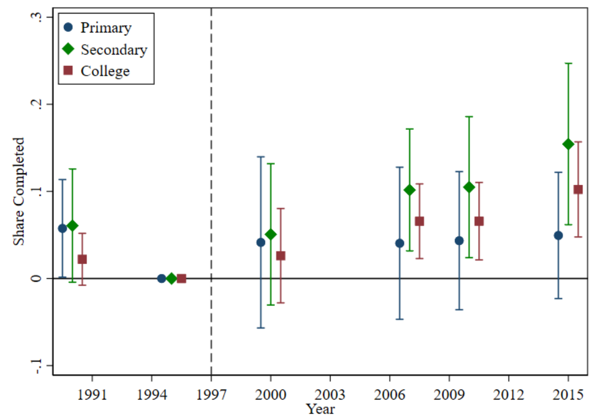

(b) Educational Attainment

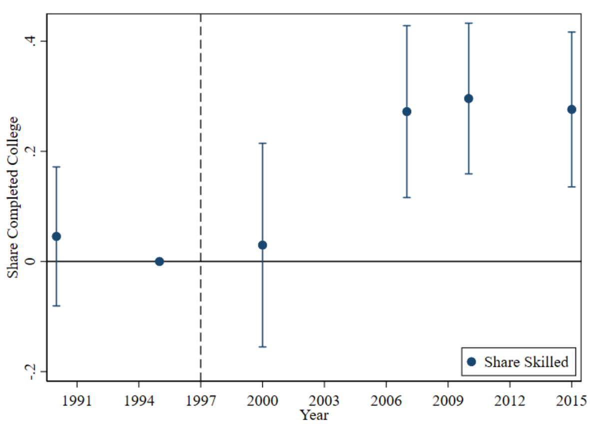

(c) Share of OFWs Skilled

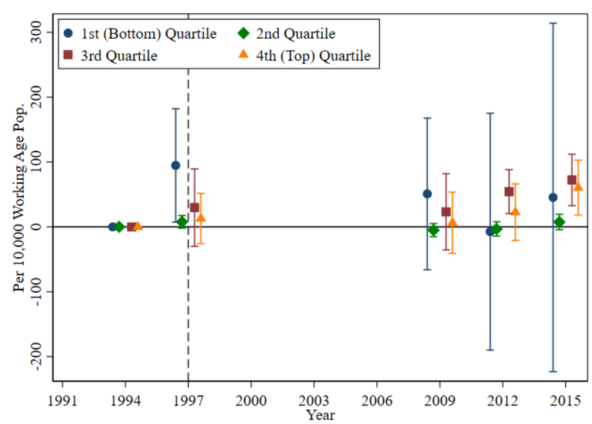

(d) OFW Occupations by Education Quartile

Note: Regressions modify Equation 2.4 to include interactions between $Shiftshare_o$ and indicator variables for each pre- and post-shock year. Panel (a) corresponds to outcomes in Table 2.6, panel (b) corresponds to outcomes in Table 2.4, and panels (c) and (d) corresponds to outcomes in Table 2.5. The 1994 or 1995 interaction term, for contract/FIES or census outcomes respectively, is omitted as the reference point. Monetary outcomes are in real 2010 PhP (PhP17.8/US$ PPP). Observations are at the province-period level. We include the partially-treated year 1997 in event study samples. 95% confidence intervals shown. Standard errors are clustered at the province level.

Table B.1: Exchange Rate Shocks and Baseline Destination Characteristics

	Dependent Variable: Exchange Rate Change ($\tilde{\Delta} R_d$)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1995 GDP Per Capita	-0.001						-0.004
	(0.007)						(0.011)
Average Contract Salary		-0.009					0.144
		(0.248)					(0.347)
Share of Contracts Professional			-0.005				-0.099
			(0.188)				(0.339)
Share of Contracts Manufacturing				-0.137			-0.317
				(0.213)			(0.253)
Share of all 1995 Contracts					0.073		0.374
					(1.011)		(1.152)
1994-1996 Exchange Rate Change						0.434	0.085
						(0.435)	(0.670)
Observations	104	104	104	104	104	104	104
Dep. Var. Mean	0.406	0.406	0.406	0.406	0.406	0.406	0.406
Dep. Var. St. Dev.	0.138	0.138	0.138	0.138	0.138	0.138	0.138
Joint F-Test P-value							0.833

Note: The table reports coefficients from regressions of the exchange rate shock on baseline destination characteristics, weighting by baseline migrant income in each destination (following Borusyak et al. (2022a)). GDP per capita is in thousands 1995 USD. Average contract salary is in millions 2010 PHPs (17.8 PhP per PPP US$ in 2010). Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Table B.2: Baseline Province Characteristics and Shock Components

	Share Rural	Asset Index	Baseline Domestic Income Per Capita	Baseline Expenditure Per Capita	Baseline Primary Sector Share	Baseline Industrial Sector Share	Baseline Service Sector Share	Baseline Financial Sector Share
$Shiftshare_o$	0.241	-1.754	-26.605	-15.362	0.153	-0.088	-0.032	-0.033
	(0.351)	(1.524)	(17.422)	(13.596)	(0.265)	(0.116)	(0.136)	(0.026)
Obs.	74	74	74	74	74	74	74	74
Dep. Var. Mean	0.643	-0.636	26.173	24.368	0.567	0.121	0.299	0.013
Dep. Var. SD	0.193	1.023	8.677	7.891	0.175	0.082	0.095	0.013

Note: Table reports coefficients from regressions for each baseline province characteristic on $Shiftshare_o$, with controls for the main effects of $MigInc_{o0}$ and $Rshock_o$. Income and expenditure are in thousand 2010 PhP (17.8 PhP per PPP US$ in 2010). Service sector excludes financial services, which is examined in as separate outcome. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Table B.3: Placebo Regressions

Variables Constructed from FIES Data
Pre Period: 1985, 1988, 1991; Post Period: 1994, 1997

			Domestic Income Subcomponents		
	(1) Domestic Income Per Capita	(2) Expenditure Per Capita	(3) Wage Income	(4) Entrepreneurial and Rental Income	(5) Other Income
$Shiftshare_o$ × Post	-2.527	2.248	-2.510	-0.205	1.530
	(15.331)	(13.258)	(9.122)	(4.831)	(4.608)
Obs.	369	369	369	369	369
Dep. Var. Mean	26.962	25.372	11.585	10.843	6.313
Dep. Var. St. Dev.	10.150	8.951	6.879	3.620	2.740

Variables Constructed from Census Data
Pre Period: 1990; Post Period: 1995

	Share Aged 20-64 Completed:			
	(1) Primary School	(2) Secondary School	(3) College	(4) Share Skilled Migrants
$Shiftshare_o$ × Post	-0.058	-0.061	-0.022	-0.045
	(0.076)	(0.062)	(0.027)	(0.104)
Obs.	148	148	148	148
Dep. Var. Mean	0.734	0.383	0.112	0.301
Dep. Var. St. Dev.	0.114	0.117	0.038	0.095

Variables Constructed from Contract Data
Pre Period: 1994; Post Period: 1997

			Contracts per 10,000 Working Age People			
	(1) Global Income Per Capita	(2) Migrant Income Per Capita	(3) 1st Quartile Education	(4) 2nd Quartile Education	(5) 3rd Quartile Education	(6) 4th Quartile Education
$Shiftshare_o$ × Post	4.693	0.579	94.758	7.695	29.691	12.897
	(20.147)	(4.369)	(125.611)	(20.999)	(175.461)	(103.683)
Obs.	148	148	148	148	148	148
Dep. Var. Mean	33.485	3.893	45.449	6.000	18.254	19.579
Dep. Var. St. Dev.	13.519	2.874	41.849	7.522	24.503	23.314

Note: Table presents coefficients on $Shiftshare_o \times Post_t$ in placebo regressions with false "post" periods. For definitions of outcomes, see: Table 2.3 (global, domestic, and income; and domestic income subcomponents), Table 2.4 (education outcomes), and Table 2.5 (share skilled migrants; migrant occupation outcomes). Compared to these other tables, $Post_t$ is redefined to refer to periods no later than 1997. All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

Table B.4: Import and Export Shocks Do Not Predict the Migrant Income Shock

	Migrant Income Shocks ($Shiftshare_o$)			
	(1)	(2)	(3)	(4)
Import Shift-share	0.000742	0.00106	-0.000121	-0.000121
	(0.00173)	(0.00160)	(0.00132)	(0.00143)
Export Shift-share	0.0000892	0.000110	-0.0000904	0.000149
	(0.00133)	(0.00165)	(0.00159)	(0.00165)
$MigInc_{o0}$	0.399***	0.562***	0.533***	0.526***
	(0.00764)	(0.0538)	(0.0539)	(0.0569)
$Rshock_o$	2.961***	2.220***	2.220***	2.238***
	(0.444)	(0.343)	(0.310)	(0.316)
Obs.	74	74	74	74
Controls	None	Panel A	Panel B	Panel C

Note: Unit of observation is province. The outcome variable is the migrant income shock $Shiftshare_o$. Controls indicate the set of additional control variables included in the regression, with panels corresponding to the structure of our main effect tables. Panel A includes destination controls only. Panel B additionally includes province development status controls. Panel C additionally includes province industrial structure controls. For list of destination and provincial controls, see Table 2.3. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Table B.5: Effects of Migrant Income Shock on Internal Migration

	Census: 1990, 2000, 2010					
	Age: 25 - 64			**Age: 16 - 24**		
	(1) In Migration Rate	(2) Out Migration Rate	(3) Net Migration Rate	(4) In Migration Rate	(5) Out Migration Rate	(6) Net Migration Rate
Panel A. Destination controls only						
$Shiftshare_o \times$ Post	-0.002	-0.016	-0.013	-0.007	-0.047	-0.040
	(0.023)	(0.012)	(0.033)	(0.028)	(0.021)**	(0.046)
Panel B. Additional province development status controls						
$Shiftshare_o \times$ Post	-0.008	-0.018	-0.010	-0.011	-0.046	-0.036
	(0.018)	(0.014)	(0.028)	(0.023)	(0.019)**	(0.038)
Panel C. Additional province industrial structure controls						
$Shiftshare_o \times$ Post	-0.020	-0.019	0.001	-0.027	-0.045	-0.019
	(0.019)	(0.011)*	(0.028)	(0.022)	(0.020)**	(0.037)
Panel D. Additional import and export shift-share variables						
$Shiftshare_o \times$ Post	-0.021	-0.019	0.002	-0.027	-0.045	-0.018
	(0.018)	(0.010)*	(0.026)	(0.019)	(0.019)**	(0.034)
Obs.	207	207	207	207	207	207
Dep. Var. Mean	0.027	0.026	-0.001	0.032	0.043	0.012
Dep. Var. St. Dev.	0.020	0.013	0.017	0.024	0.021	0.029

Note: Internal migration data is from 1990, 2010, and 2010 Censuses. Due to missing internal migration data in the 1990 Census, five provinces are dropped at the recommendation of the Philippine Statistical Authority (Camarines Sur, Capiz, Cavite, Mindoro Oriental, and Zamboanga Del Sur). Dependent variables are in-migration rate (individuals reporting having moved into the province within the last five years, as share of provincial population), out-migration rate (analogously, share who moved out of the province in the last five years), and net migration rate (the out-migration rate minus the in-migration rate). For list of destination and provincial controls, see Table 2.3. All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

Table B.6: Effects of Migrant Income Shock on Manufactured Exports

	Manufactured Exports per Capita			
	Full Period		Long Run	
	(1) Levels	(2) IHS	(3) Levels	(4) IHS
Panel A. Destination controls only				
$Shiftshare_o \times$ Post	3.486	0.589	8.318	1.359
	(10.234)	(1.096)	(15.466)	(1.697)
Panel B. Additional province development status controls				
$Shiftshare_o \times$ Post	2.870	0.275	2.492	0.831
	(12.143)	(1.279)	(18.519)	(1.920)
Panel C. Additional province industrial structure controls				
$Shiftshare_o \times$ Post	1.823	0.078	-0.569	0.475
	(12.259)	(1.306)	(18.268)	(1.920)
Panel D. Additional import and export shift-share variables				
$Shiftshare_o \times$ Post	2.989	0.233	-0.795	0.477
	(12.804)	(1.358)	(17.639)	(1.965)
Obs.	888	888	370	370
Dep. Var. Mean	2.667	0.609	2.669	0.596
Dep. St. Dev.	8.640	1.155	8.745	1.153

Note: Unit of observation is the province-year. Dependent variable is total value of manufactured exports, in thousands of real 2010 Philippine pesos (PhP), divided by province population. Dependent variable winsorized at 99%. "IHS" is inverse hyperbolic sine transformation. Manufactured exports data are from Annual Survey of Philippine Business and Industry (ASPBI), Annual Survey of Establishments (ASE) and Census of Philippine Business and Industry (CPBI) (depending on year). Full period includes all years with export data available, except the year 1997 (1994, 1996, 1998, 1999, 2006, 2008, 2009, 2010, 2012, 2013, 2014, and 2015). Long run includes years 1994, 1996, 2009, 2012, and 2015. For list of destination and provincial controls, see Table 2.3. All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

Table B.7: Effects of Migrant Income Shock on Agricultural Income

	Full Period: Triennial 1985-2018				Long Run: 1994, 2009, 2012, and 2015			
	(1) Agricultural Income	(2) Agricultural Wage Income	(3) Agricultural Non-Wage Income	(4) Non Agricultural Income	(5) Agricultural Income	(6) Agricultural Wage Income	(7) Agricultural Non-Wage Income	(8) Non Agricultural Income
Panel A. Destination controls only								
$Shiftshare_o$ × Post	-2.721	-2.209	-0.512	15.692	-2.469	-0.972	-1.498	26.287
	(3.414)	(1.083)**	(2.653)	(6.512)**	(3.823)	(1.006)	(3.172)	(3.859)***
Panel B. Additional province development status controls								
$Shiftshare_o$ × Post	0.718	-1.193	1.911	12.209	2.321	0.391	1.929	16.761
	(2.610)	(1.120)	(1.753)	(9.190)	(3.043)	(1.047)	(2.339)	(5.106)***
Panel C. Additional province industrial structure controls								
$Shiftshare_o$× Post	1.326	-1.256	2.582	13.163	2.815	0.512	2.303	16.090
	(2.869)	(0.962)	(2.027)	(8.891)	(3.279)	(0.937)	(2.570)	(5.514)***
Panel D. Additional import and export shift-share variables								
$Shiftshare_o$ × Post	1.353	-1.266	2.618	13.148	2.851	0.494	2.357	15.962
	(2.844)	(1.001)	(2.074)	(9.346)	(3.189)	(0.989)	(2.461)	(6.422)**
Obs.	813	813	813	813	296	296	296	296
Dep. Var. Mean	5.024	1.583	3.442	24.861	6.410	1.621	4.789	24.289
Dep. Var. St. Dev.	3.518	1.174	3.271	11.206	3.649	1.156	3.228	11.595

Note: Unit of observation is the province-year. Data from the Family Income and Expenditure Survey (FIES). For list of destination and provincial controls, see Table 2.3. All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

Table B.8: Exchange Rates and Foreign Direct Investment to Philippines

	FDI			
	Full Period		Long Run	
	(1) Levels	(2) IHS	(3) Levels	(4) IHS
Panel A. No Controls				
$\tilde{\Delta} R_d$ X Post	-55.635	-1.327	-68.153	-1.482
	(52.915)	(3.043)	(56.414)	(2.864)
Panel B. Destination Controls				
$\tilde{\Delta} R_d$ X Post	-17.003	0.917	-30.807	0.363
	(17.100)	(1.065)	(20.109)	(1.088)
Obs.	2,288	2,288	520	520
Dep. Var. Mean	12.145	1.788	14.221	1.941
Dep. Std. Dev.	19.237	1.782	22.425	1.806

Note: Unit of observation is country-year. Countries are weighted by the baseline migrant income in each destination. FDI data are from the PSA's Foreign Investment Reports for 1996-2002 and from PSA's OpenStat platform for after 2002. Yearly FDI are in billions of real 2010 PhPs. Full period includes years from 1996 to 2018. 1997 is dropped from the analysis due to partial treatment. Long run includes years 1996, 2009, 2012, 2015, and 2018. For list of destination controls, see Table 2.3. All regressions include province and year fixed effects. Standard errors are clustered at the country level. *** p<0.01, ** p<0.05, * p<0.10.

Table B.9: Estimating θ using Poisson Pseudo-maximum Likelihood

	OLS	PPML	PPML
	Change in Migrants		
$\text{Log}(\Delta R_d)$	9.374*	3.471**	3.417**
	(5.146)	(1.720)	(1.707)
Observations	26,344	24,788	24,788
Fixed Effects	Origin x Skill	None	Origin x Skill

Note: OLS and PPML estimates of θ using the migration response to a destination shock, at the origin-destination-skill level. Standard errors clustered at the destination level. ΔR_d is the change in exchange rates across destinations d over the course of the Asian Financial Crisis. Migrant earnings and migrant flows are from the POEA/OWWA dataset. *** indicates significance at the 1% level. ** indicates significance at the 5% level * indicates significance at the 10% level.

Table B.10: Impacts on Domestic Income by Skill, Migrant Income, and Migrant Shares

	1994, 2009, 2012, and 2015			Census
	(1)	(2)	(3)	(4)
	Domestic Income Per Capita Skilled	Domestic Income Per Capita Unkilled	Migrant Income Per Migrant	Migrant Share Age 20 - 64
Panel A. Destination controls only				
$Shiftshare_o \times$ Post	56.942	13.162	98.761	0.007
	(20.917)***	(5.346)**	(118.020)	(0.012)
Panel B. Additional province development status controls				
$Shiftshare_o \times$ Post	21.287	11.266	199.601	0.013
	(14.781)	(5.826)*	(154.790)	(0.014)
Panel C. Additional province industrial structure controls				
$Shiftshare_o \times$ Post	18.488	10.729	203.489	0.013
	(17.094)	(5.635)*	(157.317)	(0.014)
Panel D. Additional import and export shift-share variables				
$Shiftshare_o \times$ Post	18.556	10.624	208.781	0.013
	(15.272)	(5.596)*	(138.551)	(0.013)
Obs.	296	296	296	444
Dep. Var. Mean	65.934	22.362	319.519	0.018
Dep. Var. St. Dev.	18.778	7.120	104.876	0.016

Note: Unit of observation is the province-year. Overseas worker rate values are from the Census and covers 1990, 1995, 2000, 2007, 2010, and 2015. Migrant income per migrant is calculated from POEA/OWWA data. Domestic income by skill are calculated from merged Family Income and Expenditure Survey (FIES) and Labor Force Survey (LFS) data, where we define a household as skilled if any working member is skilled. For list of destination and provincial controls, see Table 2.3. All regressions include province and year fixed effects. Standard errors are exposure-robust, accounting for correlation of shocks across provinces, based on estimation of shock-level regressions (Borusyak et al., 2022a). *** p<0.01, ** p<0.05, * p<0.10.

Table B.11: Overall Changes and Model-based Decomposition of Flows and Income

	Migrant Flows	Domestic Income	Migrant Income	Global Income
Mean	0.011	26.101	4.087	30.189
Std. Dev.	(0.008)	(9.405)	(2.993)	(11.340)
Impact of 1-std.-dev. shock	**0.001**	**1.750**	**0.523**	**2.272**
Increase as % of mean	11%	6.7%	12.8%	7.5%
Share of global income increase	——	77.0%	23.0%	100.0%
Model-based decomposition:				
Education channel	16.9%	22.9%	24.1%	23.2%
Exchange rate channel	29.4%	——	74.7%	17.2%
Direct wage channel	——	60.7%	——	46.7%
Explained by model	46.2%	83.6%	98.8%	87.1%

Note: The table summarizes the changes to the variables for which we decompose the overall changes and derive the changes due to the education channel component. The mean and standard deviation values are for the closest available year before the crisis (1995 for migrant flows and 1994 income). The impact of a 1 std dev shock in migrant income is the coefficient from the regressions multiplied by 0.093 (the std. dev. of the migrant income shock). Monetary units are in thousands of Philippine pesos (PhP). The bottom panel describes the contributions of each model-based decomposition.

APPENDIX C

Appendix to Chapter III

C.1 Experiment and Intervention Details

Figure C.1: STITCH Study Experimental Design

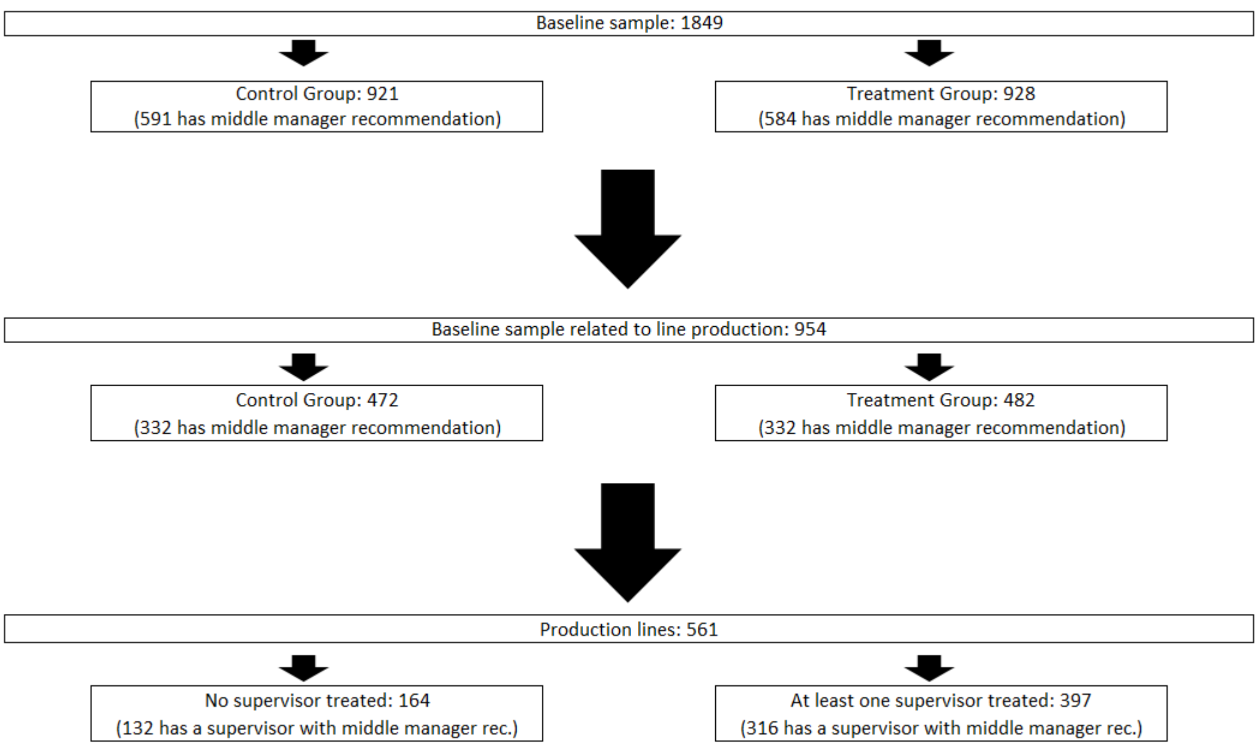

Figure C.2: Timeline of Experiment and Data Collection

Figure C.3: Distribution of Line Level Treatment, Defined as Fraction of Supervisors Treated

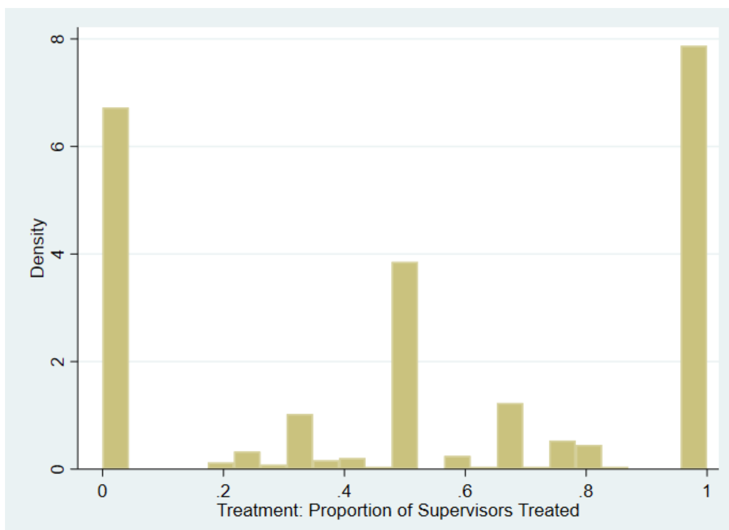

C.1.1 STITCH Modules

The STITCH training is composed of 4 modules and 25 hour-long sessions. Below, we summarize the contents of each session:

Figure C.4: Graphic Summarizing STITCH Modules and Sessions

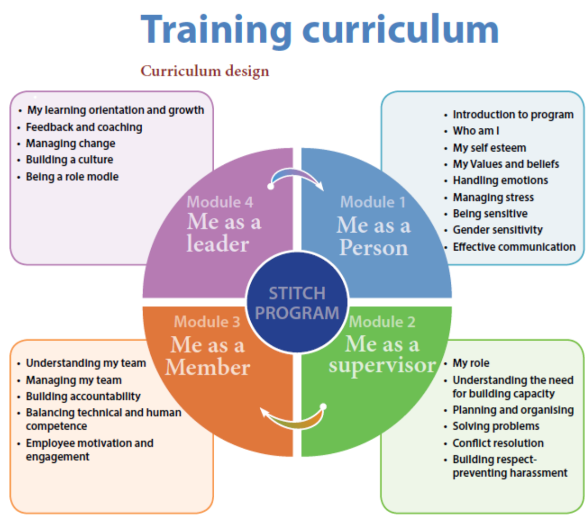

C.1.1.1 Module 1: Me as a Person

- **Introduction:** Introduction to the STITCH program. Includes interactive activities for participants to introduce themselves and to reflect on and discuss their aspirations.

- **Who am I:** Focus on enabling participants to think and reflect about themselves and what they value both in work and in life. Underlines the importance of health, knowledge/skills, and good relations for contentment in all spheres of life.

- **My Self Esteem:** Pair activity focusing on building self-esteem and confidence by better understanding of own strengths. Importance of building self-esteem within team members at work and how the fast pace of work and demand for productivity can lead to an environment detrimental to self-esteem.

- **My Behaviours and Values:** Focus on understanding behaviors required for an effective work place and the link between values and behaviors. Participants are asked to list their commitments towards developing and maintaining effective behaviors.

- **Handling Emotions:** Role-play activity to understand emotional responses and their impact. Focus on the importance of holding immediate reactions to situations and positive actions to manage emotions.

- **Managing Stress:** Role-play activity to better understand causes, reactions, and effects of stress. Tips for effective stress management, at work and in other spheres of life.

- **Being Sensitive:** Focus on the importance of sensitivity in interpersonal relationships to managing the emotions and stress of others. Role-play activity focusing on interactions between workers, supervisors, and managers to reflect on and build sensitivity.

- **Gender Sensitivity:** Focus on enabling the participants to understand the impact of socialization on attitudes and mind-sets of men and women. Reflect on how these attitudes effect behavior and outcomes at work. Participants are asked to identify one action they commit to undertaking for sensitivity.

- **Effective Communication:** Focus on understanding communication styles and develop skills to communicate assertively and responsibly. Role-playing and brainstorming activity to understand outcomes of different communication styles.

C.1.1.2 Module 2: Me as a Supervisor

- **My Role:** Focus on understanding the roles supervisors perform on a day-to-day basis. Highlight the importance of both technical responsibilities and people management responsibilities. Group discussion about what knowledge, skills, and attitudes supervisors need at possess in order to be effective in both their production and human resource management roles.

- **Understanding the Need for Building Capacity:** Focus on the importance of and tips for using time more effectively in order to find time to develop and enhance new skills. A time mapping-exercise to identify all the activities supervisors do in a given day and see how much time is spent on activity categories such as planning, problem solving, communication with management etc.

- **Planning and Organizing:** Focus on the importance of planning and organizing, especially while working as a team, and the options available within given work situations. Session composed of a team game to underline the need for planning in team work and a group activity where the groups are asked how they would plan for a hypothetical scenario (such as a specific order quantity they have to fulfill in 20 days) and discuss their options.

- **Solving Problems:** Focus on problem solving skills using case studies and role play. Underlines skills such as problem identification, analyzing the root cause of the problem, making decision based on available options, implementing the decision, and reviewing the outcomes of the decision. Includes a session where participants think of creative solutions to presented problems.

- **Conflict Resolution:** Focus on enabling supervisors to understand the causes/consequences of unresolved conflicts and helping them understand different styles of handling conflict situations (competing, collaborating, compromising, avoiding, and accommodating). Includes case studies for participants to work through and discuss.

- **Building Respect-Preventing Harassment:** Focus on helping participants understand and reflect on what constitutes harassment and its impacts, clarifying company policy regarding sexual harassment, and finding better ways to be effective at work. Case studies to clarify and discuss what constitutes harassment and how it can be prevented.

C.1.1.3 Module 3: Me as a Member

- **Understanding My Team:** Introductory session focused on the importance and benefits of teamwork. Session is mainly composed of a team game that needs to be completed without any effective communication to underline the importance of teamwork. The trainer than discusses with the participants their experiences with the activity and importance of teamwork in their work.

- **Managing My Team:** Focus on understanding various stages of team development and different leadership styles. The emphasis is on demonstrating Tuckman's stages of group development (forming, storming, norming, performing) using role playing exercises. The trainer than discusses appropriate leadership styles for the various stages.

- **Building Accountability:** Focus on the importance of responsibility and accountability for their work and functioning of their teams. At the end of the session, the participants are asked to make an action plan to create ownership and accountability within their teams.

- **Balancing Technical and Human Competence:** Focus on the importance of working on human management skills alongside technical skills to create better performing teams. Session is composed of a role-playing activity and a self-assessment questionnaire for participants to complete and reflect on.

- **Employee Motivation and Engagement:** Focus on the need for motivation/encouragement and tips for effective motivation and engagement. The session covers the Herzberg's theory of motivation and a role-playing activity to demonstrate the varying motivations of different members and how a leader can support such motivations to be effective. It emphasizes the need for showing appreciation, helping workers realize their value and encouraging learning and ownership.

C.1.1.4 Module 4: Me as a Leader

- **My Learning Orientation and Growth:** Introductory session focusing on factors that contribute to learning and growth and how to create an environment that facilitates growth. Importance of a growth-mindset as opposed to a fixed mindset.

- **Feedback and Coaching:** Focus on the importance both giving and receiving constructive feedback for developing required skills. Session is mainly composed of a role-playing activity to practice the skill of giving and eliciting both positive and negative feedback.

- **Managing Change:** Focus on the need and the impact of change and how to manage the process of change. The session includes an activity where the participants are divided into groups and each group is given a task to complete. In the middle of the process, a change is introduced to the task and participants are asked questions about the process of adaptation to this change.

- **Building a Culture:** Focus on the importance and the formation of work culture. Participants are asked to brainstorm about the traits of a good work culture and then are asked in groups to develop an action plan towards achieving a positive work culture. The trainer then highlights the importance of the supervisors for work culture and discusses how good management practices, such as providing timely feedback, can contribute to a better culture.

- **Being a Role Model:** Closing session that focuses on the concept of self-management and on developing skills and knowledge required to be a role model. The trainer explains how all components discussed in the program helps individuals develop and become role models. The participants are asked to reflect on their leanings and how these leanings can help them with their aspirations they shared in the first session.

C.2 Simple Model of Training Allocation

Our empirical results suggest that middle managers target training to supervisors who gain below the average in terms of productivity but above the average in terms of retention. In this section, we posit a simple model of supervisor training allocation to rationalize and interpret these patterns in the context of a principle agent problem. The model consists of a firm, a middle manager, and the supervisors to whom training will be allocated. Training heterogeneously affects both the individual productivity and the retention probabilities of supervisors. We study how the firm (the principle) and the middle manager (the agent) would choose to allocate the training to supervisors with heterogeneous gains.

C.2.1 Setup

Supervisors. There are two periods. A population of line supervisors are present in period 1. They have identical per-period productivity p and have quitting probability $(1-\delta)$ from period 1 to period 2. Training affects both the productivity and the quitting probability of supervisors. Supervisors are heterogeneous with regards to their responsiveness to training, where supervisor i has productivity response τ_p^i and retention response τ_δ^i. Supervisor training takes place in the first period and its productivity effects are immediately realized in period 1. So, a trained supervisor produces $p+\tau_p^i$ both periods, conditional on not quitting. The quitting probability of trained supervisors are $(1-\delta-\tau_\delta^i)$. Line supervisors who quit after period 1 are replaced. The replacement supervisors have productivity zp where $z \in (0,1]$. This term reflects that replacement supervisors can initially be less productive or that during the process of replacement the line may be less productive for a period. The replacement supervisors are not trained and their productivity are not shifted by τ_p.

Payoff for the Firm and the Middle Manager. Both the firm and the middle manager are risk neutral. The firm's objective function is to maximize line productivity.[1] The middle manager also aim to maximize line productivity, but they also incur an additional personal cost c in period 2 if the line supervisor quits after period 1. This term captures both the personal replacement/training costs the middle manager incurs, and also the fact that, in our context, part of the middle manager's job is to ensure retention of supervisors. Therefore, there can be a professional cost to high turnover of supervisors. The economic consequence is that the personal replacement cost c misaligns the principal's (the firm) and the agent's (the middle manager) objectives.

[1]Conversations with the firm confirmed that, in our context, this assumption is realistic as productivity is the largest determinant of profitability which the firm feels it can influence.

We write the middle manager's valuation of a trained and untrained supervisor i as:

$$\text{Not Trained:}\quad \underbrace{p}_{\text{Period 1 payoffs}} + \underbrace{\delta p}_{\text{Period 2 payoffs if supervisor stays}} + \underbrace{(1-\delta)(zp-c)}_{\text{Period 2 payoffs if supervisor replaced}}$$

$$\text{Trained:}\quad \underbrace{(p+\tau_p^i)}_{\text{Period 1 payoffs}} + \underbrace{(\delta+\tau_\delta^i)(p+\tau_p^i)}_{\text{Period 2 payoffs if supervisor stays}} + \underbrace{(1-\delta-\tau_\delta^i)(zp-c)}_{\text{Period 2 payoffs if supervisor replaced}}$$

The difference between the trained and the untrained value yields the middle manager payoff from training i, denoted Δ^i:

$$\Delta^i(\tau_p^i, \tau_\delta^i) = \underbrace{\tau_p^i(1+\delta) + \tau_\delta^i((1-z)p) + \tau_\delta^i\tau_p^i}_{\text{Productivity Gains} \equiv \Delta_p^i} + \underbrace{\tau_\delta^i c}_{\substack{\text{Middle Manager} \\ \text{Personal Costs} \equiv \Delta_c^i}}$$

where Δ_p^i is the fraction of the gains due to line productivity effects and Δ_c^i is the fraction due to avoiding personal replacement costs through retention effects. Given the firm's objective would be to maximize Δ_p^i when allocating the training, Δ_c^i represents the wedge in payoffs induced by middle managers personal cost to losing a supervisor. Overall line-level productivity gains from training Δ_p^i is not equal to the per-period productivity gain of the supervisor τ_p^i, as effects on retention do indirectly influence productivity through changing the probability of holding on to incumbent (and possibly more productive) supervisors and avoiding supervisor replacement costs that effect productivity (all of which are captured by the term z).

Middle Manager Information and the Distribution of Types. Supervisor types are indexed by their productivity and retention gains. The marginal distribution of the productivity gains is $\tau_p^i \sim G(.)$. We assume the middle manager has perfect knowledge about both the τ_p^i of each supervisor and the conditional average of retention gains $\mathbb{E}[\tau_\delta^i|\tau_p^i] = f(\tau_p^i)$. Critically, we assume that τ_p^i and τ_δ^i are negatively correlated ($f'(\tau_p^i) < 0$). This negative relationship induces a trade off between gains in terms of retention and productivity, and it is consistent with what we observe in the empirical analysis. Finally, we further impose $f''(\tau_p^i) < 0$ and that the ultimate retention probability $\delta + \tau_\delta$ must lie in the unit interval.

To explore why a negative relationship may exist between retention and productivity gains, suppose that the supervisor's retention response to training has two components: a component that responds to the supervisor's productivity gain from training and an idiosyncratic component. The former component can be negatively related to productivity gains, as increasing the supervisor's productivity would make them more valuable in the labor market and (without a proportionate increase in wages, which we do not observe) could increase the likelihood they leave for another job. The idiosyncratic component could be arbitrarily related to the productivity gains. In section 3.5.1.1 we show suggestive evidence that control supervisors with high recommendation had a lower retention rate compared to their low recommendation counterparts. Supposing lower baseline retention rate implies a higher retention response to training, this suggests even the component of the retention gains that is not a direct response to the changing labor market outcomes (as the control supervisors were not trained) could still be negatively correlated with productivity gains.

C.2.2 Ideal Supervisor Type to Train

We consider the question of what supervisor type has the highest training payoffs from the perspective of the middle manager, which we call the *ideal type* from the perspective of the middle manager.[2] The middle manager aims to maximize the expected payoff from training a type with τ_p^i:

$$\max_{\tau_p^i} \; \mathbb{E}[\Delta^i(\tau_p^i, \tau_\delta^i)|\tau_p^i] = \max_{\tau_p^i} \; \tau_p^i(1+\delta) + f(\tau_p^i)((1-z)p) + f(\tau_p^i)\tau_p^i + f(\tau_p^i)c$$

which leads to the ideal type:

$$\tau_p^* = \frac{1+\delta+f(\tau_p^*)}{-f'(\tau_p^*)} - p(1-z) - c. \tag{C.1}$$

and the expected productivity gain for the ideal type:

$$\mathbb{E}[\Delta_p^*] = \tau_p^*(1+\delta) + f(\tau_p^*)((1-z)p) + f(\tau_p^*)\tau_p^* \tag{C.2}$$

Remark (derived in Appendix C.2.4). The expected productivity gain from the ideal type decreases as the middle manager personal cost increases, i.e. $\frac{d\mathbb{E}[\Delta_p^*]}{dc} < 0$.

The intuition is straightforward. As the personal cost of losing a supervisor increases, the middle manager put more and more emphasis on targeting supervisors who show large

[2] Note that this is a distinct exercise from choosing to allocate the training to all the members of a type, as this would also depend on the relative density of the type in the distribution of supervisors and availability of training. Here we are simply interested in training which type of supervisor individually provides the highest value to the middle manager.

retention effects, giving up productivity gains in the process. This implies that the relative ordering of supervisors (and the allocation of scarce training) increasingly differs between the firm and the middle manager as c increases, due to the trade-off between τ_p^i and τ_c^i.

C.2.3 Treatment Effects in the Model

The Average Treatment Effect. If the firm randomizes allocation (or trains every supervisor), the average treatment effect in terms of productivity would be $\mathbb{E}[\Delta_p^i] = \mathbb{E}[\tau_p^i](1+\delta) + \mathbb{E}[f(\tau_p^i)]((1-z)p) + \mathbb{E}[f(\tau_p^i)\tau_p^i]$. This expression is the theoretical counterpart of the productivity ATE estimates we get in our empirical work as we do our analysis at the level of the production line.[3] Again, the overall productivity gains $\mathbb{E}[\Delta_p^i]$ are distinct from the supervisor level per-period productivity gains τ_p^i. By focusing on the line productivity over time irrespective of supervisor retention, our ITT estimates take into account any possible productivity gains/losses induced by changes in supervisor retention, captured in the model by retention effects and the parameter z.

Our model does not assume that the per-period treatment effects τ_p^i and τ_δ^i are distributed such that the ATE is positive. $\mathbb{E}[\Delta_p^i] > 0$ only if $G(.)$ and the $f(\tau_p)$ are such that enough weight is put on supervisors that provide an overall productivity gain to the firm.

Heterogeneity with Middle Manager Recommendation. If the firm knew $(\tau_p^i, \tau_\delta^i)$ for every supervisor, it could allocate the scarce training by ordering the supervisors by Δ_p^i and allocate the resource accordingly.[4] This would guarantee a larger treatment effect than random allocation. However, if the firm relies on middle managers to allocate the training due to information frictions, whether the treatment effects are higher or lower than random allocation depends on the relative size of the personal replacement cost c. If c is negligible, the middle manager allocation would approximately be the same as the firm allocation as the middle manager also wants to maximize line productivity. If, on the other hand, c is relatively large, middle managers could heavily target retention, generating productivity gains well below that of random assignment in the process.

[3] In the stylized model, there are no dynamics to the treatment effect and the training is instantaneous. In our empirical work, we differentiate the effects of the training while the training is ongoing and the 6-month period following the training.

[4] This decision can include not to train supervisors who would not gain from training (or would gain less than the cost of the training if there is a cost to training an individual).

C.2.4 Derivation of Remark

Remark. The expected productivity gain from the ideal type decreases as the middle manager personal cost increases, i.e. $\frac{d\mathbb{E}[\Delta_p^*]}{dc} < 0$.

Before we establish the result, we show the following lemma holds:
Lemma: $\frac{d\tau_p^*}{dc} < 0$, i.e. the productivity gain of the ideal type is decreasing in personal cost c.

Proof: Taking the total derivative of the ideal type equation 2 and reorganizing the terms, we get:

$$\frac{d\tau_p^*}{dc} = -\left(2 - \frac{f''(\tau_p)(1+\delta+\tau_p)}{f'(\tau_p)^2}\right)^{-1}$$

By assumption, $f''(\tau_p) < 0$. We also assume that supervisor gains are distributed such that $\delta + f(\tau_p) \in [0,1]$ since this expression is a probability (probability that a treated supervisor is retained for period 2). As the denominator of the last term is positive, we conclude $\frac{d\tau_p^*}{dc} < 0$. Intuitively, as middle managers incur a higher personal cost from losing supervisors, they shift the training to supervisors with relatively higher retention gains and lower productivity gains.

To show that the remark holds, we take the total derivative of the line-level expected productivity gains equation 3 with regards to personal cost c:

$$\frac{d\Delta_p^*}{dc} = \frac{d\tau_p^*}{dc}(f'(\tau_p)(p(1-z)+\tau_p^*) + 1 + \delta + f(\tau_p))$$

Plugging in the ideal type expression from equation 2 for τ_p^*, expression simplifies to:

$$\frac{d\Delta_p^*}{dc} = -\frac{d\tau_p^*}{dc} f'(\tau_p)c < 0$$

The last inequality follows from *lemma 1* and (by assumption) $f'(\tau_p) < 0$.

C.3 Decomposing Middle Manager Selection

The following framework we use closely follows Dal Bó et al. (2021).[5] We apply the model to both productivity and retention outcomes separately. Therefore, in the general framework, gains can refer either to productivity and retention.

Suppose the middle managers are perfectly knowledgeable about the gains of the supervisors from training, and they recommend supervisors both based on the gains from training (Δ_p^i) and for other idiosyncratic reasons (Δ_c^i). For example, when focusing on productivity gains from training, the idiosyncratic component can include retention gains above and beyond productivity. Denote the value of recommending supervisor i as $\tilde{\Delta}^i$:

$$\tilde{\Delta}^i = \underbrace{\beta' X_i + \eta_i}_{\equiv \tilde{\Delta}_p^i} + \underbrace{\psi' X_i + \theta_i}_{\equiv \tilde{\Delta}_c^i}$$

where X_i is the observable characteristics of supervisor i, $\tilde{\Delta}_p^i = \beta' X_i + \eta_i$ is the gains from training for supervisor i, and $\tilde{\Delta}_c^i = \psi' X_i + \theta_i$ is the idiosyncratic middle manager preferences for recommending i. Both the gains from training and the middle manager preferences have a component that can be explained by observable characteristics (β and ψ) and a component that is unobservable to the analyst (η_i and θ_i). We pool the observable and the unobservable terms together as $\Gamma \equiv \beta + \psi$ and $u_i \equiv \eta_i + \theta_i$. We then model the decision to recommend a supervisor as recommending the supervisors above a threshold (normalized to 0): $Rec_i = \mathbb{1}[\Gamma_i + u_i > 0]$. We impose further structure to the model by assuming that (η_i, θ_i) are jointly normally distributed with mean 0. This structure yields the following expected gain (derived in Appendix C.3.1):

$$\mathbb{E}[\tilde{\Delta}_p^i | X_i, Rec_i] = \beta' X_i + \rho_{u\eta}\sigma_\eta \times \lambda(X_i, Rec_i) \tag{C.3}$$

where $\lambda(X_i, Rec_i) \equiv \frac{\phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}{Rec_i - \Phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}$ is the Inverse Mill Ratio (IMR) and the $\rho_{u\eta}$ is the correlation coefficient between u_i and η_i. The IMR can be estimated using a probit regression and be plugged in as a covariate to the estimating equations.[6] If the unobservable component of gains from treatment is negatively correlated with the entire unobserved component of the middle manager selection decision (i.e. $\rho_{u\eta} < 0$), the coefficient on the IMR interac-

[5] Our setup differs from theirs in one key dimension. In their model what leads to a null relationship between the agent's selection and productivity gains is information frictions, where the agent only has imperfect information about productivity gains. If the agent's signal is very weak, agent's selection may not be related to treatment gains. We do not focus on information frictions, but instead allow for productivity related and unrelated unobservables to be negatively correlated.

[6] Specifically, parameters Γ can be estimated using a probit regression and be plugged into the IMR equation $\lambda(X_i, Rec_i)$.

tion will be negative. This would imply the unobserved component of the middle manager selection is negatively related to treatment gains. β captures the effects of all the observable components. The estimated model can be used to compare different training allocation schemes.

We employ this general framework to decompose the relationship between middle manager selection and treatment gains to observable and unobservable components, both for gains in terms of productivity (where the middle manager selection negatively predicts treatment gains) and retention (where the middle manager selection positively predicts treatment gains). Because our productivity analysis is at the level of production lines, we perform this decomposition for productivity at the line level as well, where we use the averages of observable characteristics of supervisors tied to specific lines. Further, for simplicity, we collapse the during-training and post-training treatment effects into a single after-treatment-start treatment effect for the lines.

For each outcome, the analysis proceeds in two steps. First, we use a probit model to regress middle manager high-recommendation indicator on a set of observable characteristics. We use the estimates from the first stage to calculate the inverse mill ratio (IMR), denoted $\lambda(X_l, Rec_l)$ for lines and $\lambda(X_i, Rec_i)$ for individual supervisors . We then plug in the inverse mills ratio in the following modified versions of estimating equations 3.1 and 3.2, corresponding to our productivity and retention specifications:

$$\begin{aligned} \text{Productivity:} \quad & y_{ltr} = \alpha'(\mathbb{1}[After]_{lt}\, T_l\, X_l) + \rho_{u\eta}\, \sigma_\eta\, \lambda(X_l, Rec_l) \qquad \text{(C.4)} \\ & \mathbb{1}[After]_{lt}\, T_l + \delta_l + \mu_t + \gamma_r + \epsilon_{ltr} \\ \text{Retention:} \quad & q_{tsi} = h_{0t}\, exp\,(\beta'(T_i\, X_i) + \rho_{u\eta}\sigma_\eta\, \lambda(X_i, Rec_i)) \qquad \text{(C.5)} \end{aligned}$$

where X_i and X_l are observable baseline characteristics and all lower-level interactions are included. The coefficients on the IMR identifies $\rho_{u\eta}\sigma_\eta$ for each outcome. This term has the same sign as $\rho_{u\eta}$, the correlation between unobserved component of treatment gains (η) and the total unobserved component of middle manager selection (u). In words, if the unobserved component of the middle manager selection is negatively related to treatment gains, the coefficient on the IMR should be negative.[7]

First columns of Appendix Tables C.1 and C.2 presents the probit results from the first stage. For the probit model, along with the 11 variables selected in the LASSO analysis, we include additional demographic variables (age, gender, education, local language proficiency), middle manager assessment of management skills (industrial engineering skills,

[7]For retention, a negative treatment effect implies a lower probability of quitting, so a negative coefficient on the IMR implies middle manager selection is positively correlated with retention gains.

technical skills, and management skills), and baseline management style indices from the baseline survey (initiating structure, consideration, active personnel management, and problem index). Despite the rich set of covariates included in the model, the pseudo-R^2 from this first stage is around 19.7% at the line level (for production results) and 8% at the individual level. This is consistent with our earlier conclusion that observables explain only a small fraction of the overall middle manager selection patterns.

However, the fact that we cannot explain the majority of the variation in middle manager recommendations does not necessarily mean that we cannot explain the negative relationship between the recommendation and the treatment effect of training. Observable components of the recommendation decision could still be driving the heterogeneity in treatment effects. To assess this, we turn to the second stage. The second column of Table C.1 shows that even after controlling for a rich set of controls, the coefficient on the inverse mill ratio (interacted with treatment and an indicator for after training start) is -2.3. While this coefficient is not precisely estimated with all the covariates, it is nevertheless large and indicates that the unobservable elements of middle manager recommendation are partially driving the heterogeneity in treatment effects. For retention, we get a corresponding coefficient of -0.06, which implies that the unobservable elements of middle manager recommendations are positively correlated with retention gains (keeping in mind that in the retention model a negative coefficient implies lower quitting rates, hence a higher retention).

Note on Spillovers. The proposed framework to decompose the middle manager selection into observed and unobserved components in Section 3.5.1.3 does not take potential cross-line spillovers into account. We proceed with this simplification for two reasons. First, our research question is about ascertaining whether the middle managers identify supervisors who gain the most from training and what factors drive this selection. It is not about optimal roll out scale of the training. While the inclusion of spillovers in our framework would have strong implications about the optimal scale, it is less pertinent to the question at hand. Second, our experimental design asks middle managers (who are generally in charge of one floor) to rank their supervisors based on who would gain the most from training. Therefore, in line with our research question, the middle manager decision is not about choosing the optimal training scale within the floor, but instead to identify lines/supervisors who would gain the most from training. To confirm that controlling for floor-level saturation do not change the key result that highly recommended supervisors gain less in terms of productivity, in Appendix Table C.16, we run the middle manager recommendation heterogeneity specification presented in Section 3.4.2.1, with additional controls for event period and saturation tercile interactions. Consistent with our main results, we still see that lines with highly recommended supervisors gain less from training.

C.3.1 Derivation of Equation C.3

Equation C.3 is as follows:

$$\mathbb{E}[\tilde{\Delta}_p^i|X_i, Rec_i] = \beta' X_i + \rho_{u\eta}\sigma_\eta \times \lambda(X_i, Rec_i)$$

The equation follows from standard results on multivariate normal distributions. For recommended supervisors we know the expected productivity gain from training is:

$$\mathbb{E}[\tilde{\Delta}_p^i|X_i, Rec_i = 1] = \beta' X_i + \mathbb{E}[\eta_i|u_i > -\Gamma_i]$$

Using the properties of the normal distribution and that η_i and θ_i (and consequently u_i) are mean 0, we know $\mathbb{E}[\frac{\eta_i}{\sigma_\eta}|u_i = u] = \frac{\rho}{\sigma_u}u$. Combining this with the property $\mathbb{E}[\frac{u_i}{\sigma_u}|\frac{u_i}{\sigma_u} > \frac{-\Gamma_i}{\sigma_u}] = \frac{\phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}{1-\Phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}$, we obtain:

$$\mathbb{E}[\eta_i|u_i > -\Gamma_i] = \rho\sigma_\eta \frac{\phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}{1-\Phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}$$

For non-recommended supervisors, we obtain the parallel result:

$$\mathbb{E}[\eta_i|u_i < -\Gamma_i] = \rho\sigma_\eta \frac{\phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}{-\Phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}$$

Combining these two cases yields the desired result:

$$\mathbb{E}[\tilde{\Delta}_p^i|X_i, Rec_i] = \beta' X_i + \rho\sigma_\eta \frac{\phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)}{Rec_i - \Phi\left(\frac{-\Gamma_i' X_i}{\sigma_u}\right)} \equiv \beta' X_i + \rho\sigma_\eta \lambda(X_i, Rec_i)$$

Table C.1: Selection and Production Effect Heterogeneity

	First Stage **1[High Rec]** (1)	Second Stage **Efficiency** (2)
Treatment (ATE)		2.769**
		(1.298)
Inverse Mills Ratio		-2.309
		(1.479)
Age	-0.039*	0.716**
	(0.023)	(0.294)
1(Male)	-0.271	2.944
	(0.255)	(3.406)
1(Finished Highschool)	0.182	-11.933*
	(0.439)	(6.514)
Local Language Proficiency	-0.242	-4.062
	(0.221)	(3.017)
Supervised Dif Line Before	0.626***	-4.405
	(0.224)	(2.763)
Ever Worked as Operator	0.493	-2.592
	(0.317)	(4.220)
Ever Worked at Another Factory	0.025	3.048
	(0.236)	(2.979)
Months as Supervisor	-0.001	0.028
	(0.003)	(0.043)
Months Supervising Current Line	0.001	-0.053
	(0.005)	(0.070)
Years in Shahi	0.026	-0.334
	(0.024)	(0.340)
Motivation to Imporove (Scored by Middle Manager)	0.243*	0.574
	(0.129)	(1.921)
Months as Supervisor (Answered by Middle Manager)	-0.001	-0.108*
	(0.004)	(0.057)
Target Effort Index	0.379***	-5.551***
	(0.124)	(1.609)
Cognitive Ability	-1.212**	-7.059
	(0.519)	(6.955)
Technical Skills (Scored by Middle Manager)	-0.041	3.136
	(0.129)	(2.188)
Industrial Engineering Skills (Scored by Middle Manager)	-0.180	-1.521
	(0.142)	(2.065)
Management Skills (Scored by Middle Manager)	0.010	-2.685
	(0.138)	(2.166)
Self Esteem	0.653***	4.602*
	(0.206)	(2.775)
Initiating Structure	0.016	-0.394
	(0.023)	(0.403)
Consideration	-0.039	0.843**
	(0.032)	(0.399)
Active Personnel Management	0.018	2.236
	(0.118)	(1.554)
Problem Index	-0.165	0.789
	(0.114)	(1.372)
Baseline Productivity of Line	-0.005	0.120
	(0.005)	(0.080)
1(From Different State)	0.520*	-4.722
	(0.308)	(4.684)
1(General Caste)	-0.297	1.958
	(0.221)	(3.095)
Observations	379	
Pseudo R-sq	0.197	

Note: *** p<0.01, ** p<0.05, * p<0.10. The first column shows the results from the first stage probit model for line level middle manager selection. Second column presents results from the heterogeneous treatment effect regression. For column 2, all shown coefficients are for the triple interaction of variable of interest with treatment and 1[After Training Start].

Table C.2: Selection and Retention Effect Heterogeneity

	First Stage **Middle Manager Selection** (1)	Second Stage **1[Quit]** (2)
Treatment		-0.106
		(0.092)
Inverse Mills Ratio		-0.058
		(0.112)
Age	-0.001	0.022
	(0.008)	(0.017)
1(Male)	-0.120	0.250
	(0.108)	(0.231)
1(Finished Highschool)	0.186	0.107
	(0.172)	(0.350)
Local Language Proficiency	0.085	0.217
	(0.087)	(0.196)
Supervised Dif Line Before	0.199**	-0.408
	(0.094)	(0.461)
Ever Worked as Operator	0.507***	-0.455**
	(0.112)	(0.200)
Ever Worked at Another Factory	0.260**	0.166
	(0.107)	(0.247)
Months as Supervisor	-0.000	0.005**
	(0.001)	(0.002)
Months Supervising Current Line	-0.001	-0.010***
	(0.001)	(0.003)
Years in Shahi	-0.003	0.020
	(0.011)	(0.028)
Motivation to Improve (Scored by Middle Manager)	0.273***	-0.165
	(0.065)	(0.128)
Months as Supervisor (Answered by Middle Manager)	0.004***	0.004
	(0.002)	(0.005)
Target Effort Index	0.005	-0.244***
	(0.051)	(0.072)
Cognitive Ability	-0.374	-0.312
	(0.239)	(0.641)
Technical Skills (Scored by Middle Manager)	-0.159**	0.181
	(0.065)	(0.198)
Industrial Engineering Skills (Scored Middle Manager)	-0.151**	0.073
	(0.067)	(0.161)
Management Skills (Scored by Middle Manager)	0.084	0.261
	(0.066)	(0.161)
Self Esteem	0.122	-0.203
	(0.086)	(0.305)
Initiating Structure	0.024**	0.040*
	(0.011)	(0.023)
Consideration	-0.028**	-0.022
	(0.013)	(0.032)
Active Personnel Management	-0.008	-0.016
	(0.048)	(0.086)
Problem Index	-0.018	0.123
	(0.046)	(0.092)
1(From Different State)	0.110	0.382
	(0.130)	(0.300)
1(General Caste)	0.001	0.433**
	(0.093)	(0.176)
Observations	867	866
Pseudo R-sq	0.078	

Note: *** p<0.01, ** p<0.05, * p<0.10. The first column shows the results from the first stage probit model for supervisor level middle manager selection. Second column presents results from the heterogeneous treatment effect cox regression for retention. For column 2, all shown coefficients are for the interaction of variable of interest with treatment indicator.

Figure C.5: Random Allocation vs. Middle Manager Allocation

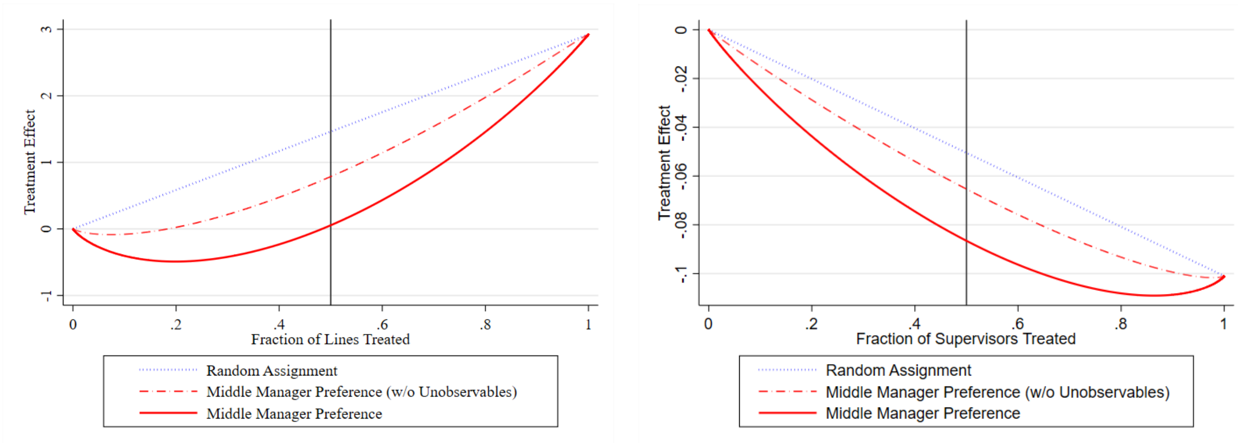

(a) Line Productivity

(b) Supervisor Retention

Note: Treatment effects with random allocation and middle manager allocation. Left panel shows results for line productivity. Right panel shows results for supervisor retention. Horizontal line signifies the point where 50% of lines (for productions) or supervisors (for retention) are treated.

Figure C.6: Training Allocation Based on Middle Manager Skill Scores

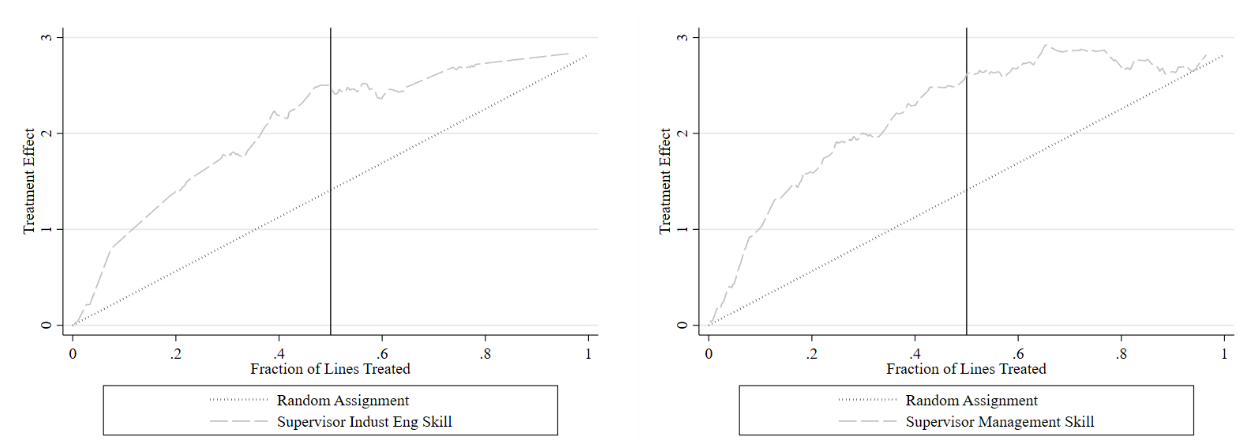

(a) Allocate Based on Industrial Engineering Skill Score

(b) Allocate Based on Management Skill Score

Note: Productivity effects with allocation rules based on middle manager assessment of supervisors. Left panel allocates based on the industrial engineering score. Right panel allocates based on the management skill score. Training is allocated first to lines with lowest average supervisor skill scores. Horizontal line signifies the point where 50% of lines are treated.

C.4 Analysis Details

C.4.1 Estimation of Baseline Line Productivity

Following the methodology outlined in Adhvaryu et al. (2022d), we estimate the baseline line productivity using a two-way fixed effect model that matches garment styles, days, and production lines. This methodology is parallel to the worker-firm mathching model of Abowd et al. (1999). We project line-day level productive efficiency on line, day, and garment style fixed effects.[8] We do this analysis for January 2007 to March 2007, the three months preceding the beginning of training. We use the fixed effect estimates for each line as the baseline productivity of the line.

C.4.2 Lasso Procedure and Included Variables

We use Stata's *lasso* command to implement a linear lasso model where the penalty term λ is selected through 10-fold cross validation. The selected $\lambda = 0.094$ with an out of sample R^2 of 2.4%. The list of variables included in the lasso procedure is below. Appendix section C.4.3 provide further details on the how the personality and management style indices below are created from our surveys.

- **Demographics, Tenure, and Experience:** Age (with age squared), Gender, $\mathbb{1}$[Finished high school], $\mathbb{1}$[General caste], $\mathbb{1}$[From out of state], $\mathbb{1}$[Native language is local language], $\mathbb{1}$[Hindu], Tenure in garment industry (months), tenure as supervisor (months), months supervising current line, tenure in Shahi (years), Ever worked as operator, Supervised different line before, Worked at different factory before
- **Middle Manager and Supervisor Joint Characteristics**: From same state, Same religion, Same gender, Supervisor hired after middle manager, Coincident tenure
- **Personality:** Conscientiousness, Locus of Control, Perseverance, Self Esteem
- **Management Style and Practices:** Consideration, Initiating structure, Conflict Index, Problem Index, Autonomous problem solving, Target effort index, Monitoring frequency, Communication index, Active personnel management
- **Self Assessment:** Technical tailoring skill, Industrial engineering skill, Managerial skill, Training interest, Expected gain from training, Amount supervisor would allocate to training, Self efficacy index, Instrumentality of training

[8]Our data includes the garment style a line produces in a given day.

- **Moddle Manager Assessment of Supervisor:** Technical tailoring skill, Industrial engineering skill, Managerial skill, Motivation to improve, Months supervising current line
- **Other:** Cognitive ability, Risk preference, Discount index, Baseline line productivity, Suggested hires last month

C.4.3 Creation of Survey Indices

The table below outlines the details of the questions used and how they are combined for the creation of the indices from the baseline supervisor surveys.

C.5 Additional Checks and Results

C.5.1 Additional Line Balance

Table C.3: Line Level Descriptive Statistics and Balance for Analysis Subsets

	Analysis Subsample				Analysis Subsample w/ Middle Manager			
	Num. Lines	Mean	SD	Coefficient (SE)	Num. Lines	Mean	SD	Coefficient (SE)
Baseline Productive Efficiency	476	55.88	13.01	-3.143**	393	55.71	12.80	-1.625
				(1.513)				(1.695)
Baseline Attendance	471	0.90	0.05	-0.009	393	0.90	0.05	-0.011*
				(0.006)				(0.006)
Baseline Retention	465	0.84	0.13	0.010	389	0.84	0.12	-0.013
				(0.016)				(0.016)
Baseline SAM	476	55.79	20.41	-1.638	393	55.71	20.27	2.268
				(2.462)				(2.829)
Baseline Budgeted Efficiency	476	60.97	7.17	-0.377	393	60.27	7.17	-1.468
				(0.841)				(0.966)
Baseline Number of Operators	476	171.32	171.75	-6.450	393	168.37	172.37	-21.155
				(18.931)				(21.741)

Note: The left panel excludes days with 0 efficiency and excludes lines that record over 20% of the days as 0 efficiency before, during or after training. The right panel further reduces the sample to lines we have middle manager recommendations for. The coefficient(SE) is from regressing the outcome on the continuous treatment indicator. Robust standard errors are reported (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.). All baseline values are from 3 months preceding training start (January - March 2017). Baseline (budgeted) efficiency is an average of daily (budgeted) efficiency values for this period. Baseline attendance and retention are the attendance and retention outcomes for the workers we matched to these lines using the personnel rosters.

C.5.2 Pre-Post Test Scores ANCOVA Analysis

To assess the effects of training on learning the module content, we present the results of the following ANCOVA specification:

$$s_{i2} = \beta_0 + \beta_1 T_i + \beta_2 s_{is1} + \mu_s + \epsilon_{i2} \tag{C.6}$$

where s_{i2} is the post-module test score in percentage points of supervisor i, s_{i1} is the pre-module test score, T_i is whether the supervisor is randomized into treatment, and μ_s is strata fixed effects. Results for all 4 modules are presented in Table C.4. Across all modules, treatment leads to a significant gain in the post-module tests. Appendix Table C.5 presents the null results on heterogeneous effects by middle manager recommendation.

Table C.4: Treatment Effect on Post-Module Exam Scores

	Dependent Variable: Post-Module Test Score			
	Module 1	Module 2	Module 3	Module 4
	(1)	(2)	(3)	(4)
Treatment	22.130***	22.048***	32.248***	38.799***
	(1.440)	(0.846)	(3.603)	(2.216)
Observations	623	574	553	541
Control Mean of Dependentent Variable	48.246	54.605	31.579	35.714
Strata FE	Yes	Yes	Yes	Yes

Note: Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1). Specification includes controls for the pre-module test scores. Test scores are in percentage points.

Table C.5: Treatment Effect Heterogeneity by Middle Manager Recommendation on Post-Module Exam Scores

	Dependent Variable: Post-Module Test Score			
	Module 1	Module 2	Module 3	Module 4
	(1)	(2)	(3)	(4)
Treatment	19.510***	19.268***	28.456***	33.631***
	(2.765)	(2.485)	(5.435)	(4.312)
High Rec.	-3.805	0.232	-2.794	-1.353
	(3.004)	(3.207)	(5.487)	(4.558)
Treatment $\times$ High Rec.	4.500	3.261	4.506	3.304
	(3.361)	(3.408)	(6.123)	(5.104)
Observations	389	364	352	338
Control Mean of Dependentent Variable	47.743	55.208	33.333	39.286
Strata FE	Yes	Yes	Yes	Yes

Note: Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1). Specification includes controls for the pre-module test scores. Test scores are in percentage points.

C.5.3 Dynamic Specification for Productivity Results

Appendix Figure C.7 charts the monthly coefficients estimates β_m from the following event study specification:

$$y_{ltr} = \sum_{m=-6}^{20} \beta_r (T_l \times D_{ltr}^m) + \delta_l + \mu_t + \gamma_r + \epsilon_{ltr} \tag{C.7}$$

where y_{ltr} is productive efficiency in line l, date t, and relative date r, T_l is line level treatment, D_{ltr}^m is an indicator for whether the date is within m months since treatment start for the factory, and δ_l, μ_t, γ_r are fixed effects.

Figure C.7: Event Study Results

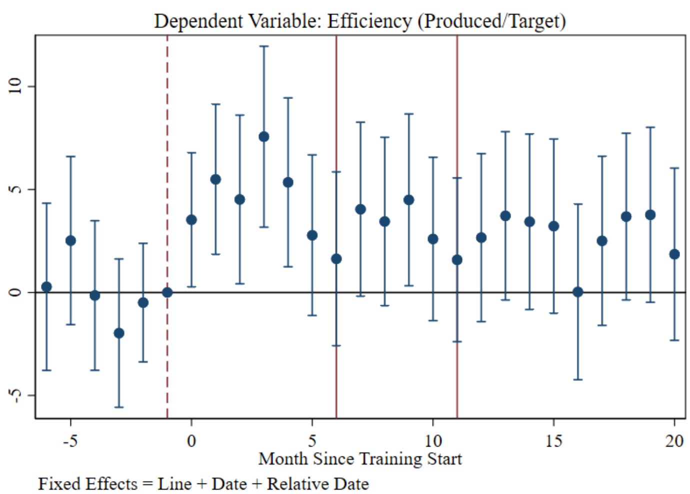

Note: Figure shows β_m from estimating equation C.7. Month 0 signifies treatment start. β_{-1} is normalized to 0. Shortest training duration is 6-months (first solid red line). Longest training duration is 11 months (second solid red line). 95% confidence intervals are shown.

C.5.4 Additional Productivity Results Tables

Table C.6: Effects of Training on Line Productivity - Panel Balanced on Relative Month

	Outcome: Efficiency (Produced/Target)				
	Analysis Lines			Lines w/ Middle Manager Match	All Lines
	(1)	(2)	(3)	(4)	(5)
During Training X Treatment	4.002***	4.009***	4.092***	3.865***	4.300***
	(1.109)	(1.111)	(1.111)	(1.247)	(1.286)
After Training X Treatment	2.789**	2.791**	2.893**	2.848**	2.829*
	(1.260)	(1.262)	(1.228)	(1.433)	(1.575)
Observations	274300	274299	274299	226915	305341
Number of Lines	480	480	480	395	553
Cont. Mean of Dep. Var.	55.533	55.533	55.533	55.533	55.533
Line FE	X	X	X	X	X
Month FE	X				
Day FE		X	X	X	X
Relative Date FE			X	X	X

Note: Standard errors are clustered at line level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). The analysis covers six months prior to twenty months after training start for each line. Days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis. Column (5) includes both the dropped lines and the line-days with 0 efficiency.

Table C.7: Productivity Effect Heterogeneity by Middle Manager Recommendation

	Analysis Lines			All Lines
	(1)	(2)	(3)	(4)
During Training X Treatment	6.023***	6.015***	6.183***	6.832***
	(1.951)	(1.955)	(1.928)	(2.056)
After Training X Treatment	6.414***	6.383**	6.617***	6.350**
	(2.466)	(2.470)	(2.431)	(2.556)
During Training X High Rec X Treatment	-4.456*	-4.428*	-4.654*	-5.946**
	(2.502)	(2.506)	(2.489)	(2.814)
After Training X High Rec X Treatment	-6.458**	-6.408**	-6.739**	-5.993*
	(3.168)	(3.172)	(3.118)	(3.523)
Observations	189381	189380	189380	208691
Number of Lines	395	395	395	444
Control Mean of Dependent Variable	55.279	55.279	55.279	55.279
Line FE	X	X	X	X
Month FE	X			
Day FE		X	X	X
Relative Date FE			X	X

Note: Standard errors are clustered at line level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). The analysis covers six months prior to training start month and the six months post the training end month for each factory. For columns (1) - (3) days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis. Column (4) includes both the dropped lines and the line-days with 0 efficiency. "High Rec" is an indicator for whether the line has average supervisor recommendation above the median.

Table C.8: Average Productivity Effects and Middle Manager Recommendation Heterogeneity - Robustness to Alternative Samples

	Dependent Variable: Efficiency (Produced/Target)					
	Average Treatment Effect			Recommendation Heterogeneity		
	All Analysis Lines	All Sup. High or Low Rec.	All Sup. Treated or Control	All Analysis Lines	All Sup. High or Low Rec.	All Sup. Treated or Control
	(1)	(2)	(3)	(4)	(5)	(6)
During Training X Treatment	4.089***	3.967***	3.806***	6.183***	6.396***	5.654***
	(1.116)	(1.354)	(1.147)	(1.928)	(2.056)	(2.011)
After Training X Treatment	3.267**	3.597**	3.115**	6.617***	7.387***	6.581**
	(1.338)	(1.683)	(1.377)	(2.431)	(2.535)	(2.578)
During Training X High Rec X Treatment				-4.654*	-4.802*	-4.299*
				(2.489)	(2.698)	(2.577)
After Training X High Rec X Treatment				-6.739**	-7.517**	-7.062**
				(3.118)	(3.316)	(3.259)
Observations	228166	151104	139940	189380	151104	107163
Number of Lines	480	314	300	395	314	227
Control Mean of Dependent Variable	55.865	55.462	55.865	55.279	55.462	55.279

Note: Standard errors are clustered at line level (*** p<0.01, ** p<0.05, * p<0.10). Specification includes day, line, and relative day fixed effects. The analysis covers six months prior to training start month and the six months post the training end month for each factory. Days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis. "High Rec" is an indicator for whether the line has average supervisor recommendation above the median. Columns (1) and (4) replicate main results. Columns (2) and (5) only includes lines where all supervisors are either high- or low-recommendation. Columns (3) and (6) only includes lines where all line supervisors are treated or control.

C.5.5 Productivity Results with Binary Treatment Definition

We replicate our line-level productivity results with an alternative binary line-level treatment definition: at least one line supervisor is treated. Given lines have varying number of supervisors, the probability of having at least one treated supervisor differs across lines. Therefore, for each line, we calculate the probability of having at least one treated supervisor given our randomization scheme. We then show results for three weighting schemes: (1) unweighted, (2) inverse probability weighted (IPW) to recover ATEs, and (3) IPW with lines with extreme weights dropped. Appendix Table C.9 shows average productivity results, corresponding to Table 3.3 in the body. Appendix Figure C.8 shows monthly event studies estimated using estimating equation C.7. Appendix Table C.10 show middle manager recommendation heterogeneity. Overall, our results are consistent across continuous and binary treatment definitions.

Figure C.8: Monthly Event Study Results - Binary Treatment

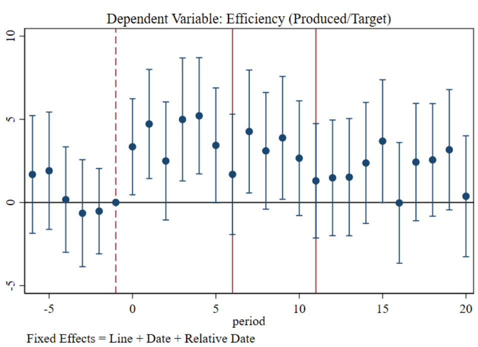

(a) Unweighted

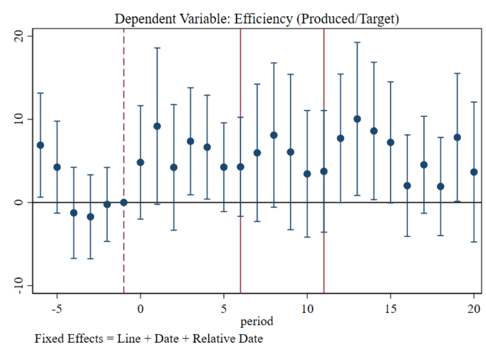

(b) IPW

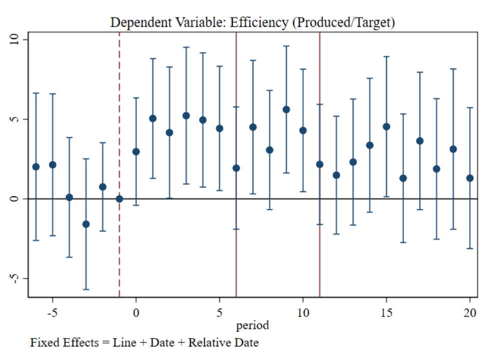

(c) IPW - Drop if Weight > 20

Note: Figure shows β_m from estimating equation C.7. Month 0 signifies treatment start. β_{-1} is normalized to 0. Shortest training duration is 6-months (first solid red line). Longest training duration is 11 months (second solid red line). 95% confidence intervals are shown.

Table C.9: Effects of Training on Line Productivity - Binary Treatment

	Outcome: Efficiency (Produced/Target)								
	No Weighting			IPW			IPW, Drop If Weight > 20		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment X During Training	3.215***	3.221***	3.095***	4.141*	4.156*	4.585**	3.517***	3.520***	3.561***
	(1.023)	(1.025)	(1.020)	(2.189)	(2.202)	(2.039)	(1.201)	(1.204)	(1.206)
Treatment X After Training	2.450**	2.439**	2.274**	5.704**	5.657**	5.993**	2.703*	2.686*	2.805**
	(1.131)	(1.133)	(1.119)	(2.622)	(2.640)	(2.614)	(1.378)	(1.383)	(1.362)
Observations	228167	228166	228166	228167	228166	228166	222350	222349	222349
Number of Lines	480	480	480	480	480	480	468	468	468
Control Mean of Dependent Variable	55.865	55.865	55.865	55.865	55.865	55.865	55.865	55.865	55.865
Line FE	X	X	X	X	X	X	X	X	X
Month FE	X			X			X		
Day FE		X	X		X	X		X	X
Relative Date FE			X			X			X

Note: Standard errors are clustered at line level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). Treatment is a binary indicator for having at least one treated supervisor on the line. The analysis covers six months prior to training start month and the six months post the training end month for each factory. For columns (1) - (4) days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis.

Table C.10: Productivity Effect Heterogeneity by Middle Manager Recommendation - Binary Treatment

	Outcome: Efficiency (Produced/Target)								
	No Weighting			IPW			IPW, Drop If Weight > 20		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment X During Training	4.320**	4.312**	4.314**	7.763**	7.755**	8.032**	6.087***	6.088***	6.068***
	(1.694)	(1.699)	(1.679)	(3.386)	(3.392)	(3.206)	(2.196)	(2.200)	(2.183)
Treatment X After Training	4.457**	4.427**	4.429**	10.843***	10.731***	11.060***	6.621***	6.598***	6.541***
	(1.975)	(1.978)	(1.965)	(4.005)	(4.036)	(4.116)	(2.226)	(2.231)	(2.218)
Treatment X During X High Rec	-3.263	-3.235	-3.548	-8.434**	-8.413**	-8.706**	-5.550**	-5.552**	-5.587**
	(2.224)	(2.229)	(2.233)	(3.809)	(3.806)	(3.610)	(2.651)	(2.652)	(2.652)
Treatment X After X High Rec	-4.364*	-4.319*	-4.718*	-11.276***	-11.140***	-11.404**	-7.636**	-7.624**	-7.632**
	(2.509)	(2.513)	(2.499)	(4.269)	(4.285)	(4.418)	(3.000)	(3.003)	(2.969)
Observations	189381	189380	189380	189381	189380	189380	183564	183563	183563
Number of Lines	395	395	395	395	395	395	383	383	383
Control Mean of Dependent Variable	55.279	55.279	55.279	55.279	55.279	55.279	55.279	55.279	55.279
Line FE	X	X	X	X	X	X	X	X	X
Month FE	X			X			X		
Day FE		X	X		X	X		X	X
Relative Date FE			X			X			X

Note: Standard errors are clustered at line level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). Treatment is a binary indicator for having at least one treated supervisor on the line. The analysis covers six months prior to training start month and the six months post the training end month for each factory. For columns (1) - (4) days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis.

C.5.6 Alternative Explanations Tables

Table C.11: Middle Manager Recommendation Is Not Well Explained by Demographic and Favoritism Related Variables

	Middle Manager Recommendation	
	All Supervisors (1)	Productivity Sample (2)
Supervisor Age	0.015	0.036
	(0.063)	(0.106)
Supervisor Age Squared	-0.000	-0.001
	(0.001)	(0.002)
Supervisor 1(Male)	-0.161	0.318
	(0.276)	(0.334)
Supervisor 1(Hindu)	0.666	0.803
	(0.442)	(0.688)
Supervisor Native Language is Kannada	-0.394***	0.102
	(0.144)	(0.229)
Supervisor from Different State	0.173	0.508***
	(0.144)	(0.190)
Supervisor 1(General Caste)	-0.084	-0.176
	(0.107)	(0.142)
Sup and Middle Manager Same Gender	0.137	-0.277
	(0.273)	(0.330)
Sup and Middle Manager Same Age Group	0.111	0.052
	(0.134)	(0.179)
Sup and Middle Manager Same Caste	0.119	0.190
	(0.106)	(0.140)
Sup and Middle Manager Same Religion	-0.368	-0.724
	(0.362)	(0.592)
Sup and Middle Manager Coincident Tenure (Years)	-0.003	0.002
	(0.015)	(0.022)
Supervisor Hired After Middle Manager	-0.005	-0.141
	(0.117)	(0.166)
Sup and Middle Manager Same Cohort	-0.163	0.031
	(0.232)	(0.349)
Constant	2.711**	2.251
	(1.067)	(1.761)
Observations	1051	585
R Sq.	0.016	0.025
F-stat	1.359	1.090

Note: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. Coefficients are from a linear regression of middle manager recommendation of a supervisor on supervisor and joint supervisor and middle manager characteristics.

Table C.12: Middle Manager Recommendation and Middle Manager Skill Scores

	Middle Manager Recommendation			
	All Supervisors		Productivity Sample	
	(1)	(2)	(3)	(4)
Technical Tailoring Skills	-0.102*	-0.145**	0.009	-0.060
	(0.060)	(0.058)	(0.087)	(0.083)
Industrial Engineering Skills	-0.058	-0.089	-0.021	-0.050
	(0.067)	(0.064)	(0.097)	(0.090)
Management Skills	0.093	-0.046	0.077	-0.110
	(0.060)	(0.063)	(0.086)	(0.087)
Motivation to Improve		0.347***		0.489***
		(0.061)		(0.078)
Constant	3.343***	2.748***	2.812***	1.904***
	(0.213)	(0.236)	(0.292)	(0.317)
Observations	1289	1289	695	695
R Sq.	0.005	0.027	0.002	0.046
F-stat	2.150	10.605	0.413	10.004

Note: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. Coefficients are from a linear regression of middle manager recommendation of a supervisor on the middle manager skill assessment scores.

Table C.13: Baseline Productivity and Middle Manager Assessment of Skills

	Dependent Variable: Baseline Productivity				
	(1)	(2)	(3)	(4)	(5)
Supervisor Technical Skills	2.667***				1.018
	(0.880)				(1.346)
Supervisor Industrial Engineering Skills		3.130***			3.737**
		(0.931)			(1.590)
Supervisor Management Skills			1.536		-1.483
			(0.935)		(1.437)
Supervisor Motivation				0.821	-0.239
				(1.061)	(1.203)
Observations	393	393	393	393	393
R Sq.	0.019	0.031	0.008	0.002	0.036
F-Statistic					3.725

Note: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. Coefficients are from a linear regression of baseline productivity on the specific middle manager assessment on skill. Middle manager assessment of supervisor skills are aggregated at the line level by taking the average of all the line supervisors. Baseline productivity is calculated as described in Appendix Section C.4.1.

C.5.7 Additional Retention Tables

Table C.14: Retention Effects are Driven By High Recommendation Supervisors

	Dependent Variable: Supervisor Quit					
	Cox Prop. Hazard			OLS		
	High Rec. (1)	Low Rec. (2)	Pooled (3)	High Rec. (4)	Low Rec. (5)	Pooled (6)
Treatment	-0.324**	0.099	0.098	-0.132**	0.034	0.029
	(0.162)	(0.148)	(0.146)	(0.051)	(0.049)	(0.048)
Treatment X High Rec.			-0.381*			-0.147**
			(0.208)			(0.069)
Observations	421	457	885	421	457	885
Relative Hazard of Treatment	0.723	1.104	-	-	-	-
Control Mean of Dep. Var.	-	-	-	0.537	0.429	0.483

Note: *** p<0.01, ** p<0.05, * p<0.10. Strata fixed effects and main effect of high recommendation dummy are included in all specifications. The outcome variable for the OLS specification is a dummy for whether the supervisor quit between treatment start and end of 2018. The sample is restricted to supervisors that could be matched to the attendance roster and supervisors who did not quit the firm between the baseline survey and the training start in their factories. Columns 1 and 4 (2 and 5) limit the sample to supervisors with high (low) middle manager recommendation. Column 3 and 6 pools both groups. Relative Hazard is calculated as the exponent of the coefficient on treatment.

C.5.8 Additional Spillover Analysis Tables and Figures

Figure C.9: Supervisor Retention by Linemate Treatment

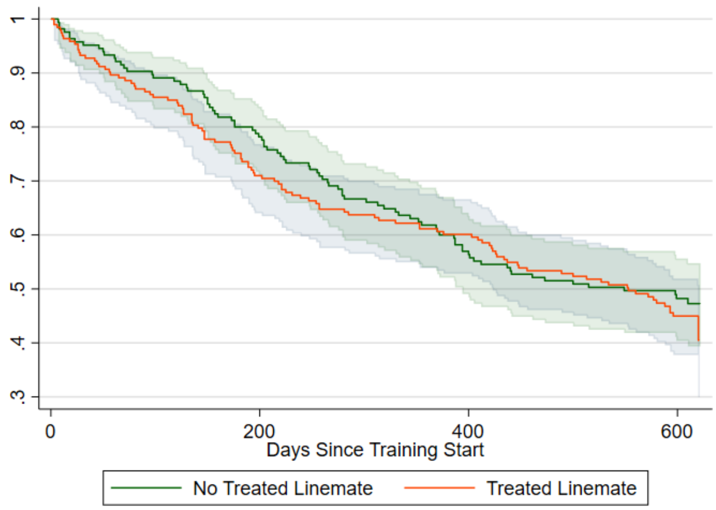

Note: Figure plots retention curves for control supervisors with and without treated supervisors on their lines. 95% confidence intervals shown.

Table C.15: Spillovers for Line Productivity Effects

	Efficiency (Produced/Target)		
	(1)	(2)	(3)
During Training X Treatment	4.476*	4.559*	3.749
	(2.555)	(2.578)	(2.543)
After Training X Treatment	0.264	0.326	-0.568
	(3.309)	(3.328)	(3.333)
During Training X Second Tercile of Saturation	1.729	1.818	1.014
	(1.706)	(1.723)	(1.715)
During Training X Third Tercile of Saturation	0.108	0.073	-0.002
	(2.304)	(2.316)	(2.268)
After Training X Second Tercile of Saturation	4.402**	4.374**	3.660*
	(1.859)	(1.868)	(1.883)
After Training X Third Tercile of Saturation	3.487	3.365	3.579
	(2.759)	(2.767)	(2.639)
During Training X Second Tercile of Saturation X Treatment	-3.254	-3.362	-2.749
	(3.340)	(3.365)	(3.304)
During Training X Third Tercile of Saturation X Treatment	0.894	0.886	1.935
	(3.442)	(3.465)	(3.393)
After Training X Second Tercile of Saturation X Treatment	-1.618	-1.643	-0.970
	(4.004)	(4.033)	(4.038)
After Training X Third Tercile of Saturation X Treatment	2.276	2.218	3.399
	(4.515)	(4.537)	(4.421)
Observations	197639	197638	197638
Number of Lines	422	422	422
Number of Floors	54	54	54
Control Mean of Dependent Variable	54.447	54.447	54.447
Line FE	X	X	X
Month FE	X		
Day FE		X	X
Relative Date FE			X

Note: Standard errors are clustered at line level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). Saturation is defined as the fraction of supervisors treated on the production floor. The analysis covers six months prior to training start month and the six months post the training end month for each factory. For columns (1) - (4) days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis.

Table C.16: Treatment Effect Heterogeneity by Middle Manager Recommendation Controlling for Spillovers

	Efficiency (Produced/Target)		
	(1)	(2)	(3)
During Training X Treatment	5.616***	5.628***	5.480***
	(2.119)	(2.131)	(2.091)
After Training X Treatment	3.787	3.730	3.628
	(2.615)	(2.625)	(2.593)
During Training X High Rec X Treatment	-4.096*	-4.024	-4.342*
	(2.481)	(2.501)	(2.450)
After Training X High Rec X Treatment	-6.140*	-5.967*	-6.447**
	(3.343)	(3.363)	(3.274)
Observations	168335	168334	168334
Number of Lines	356	356	356
Number of Floors	51	51	51
Control Mean of Dependent Variable	54.447	54.447	54.447
Line FE	X	X	X
Month FE	X		
Day FE		X	X
Relative Date FE			X

Note: Standard errors are clustered at line level (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). All specifications include controls for the interaction of floor saturation tercile and event period (during/after training). "High Rec" is an indicator for whether a line has average supervisor recommendation above the median. The analysis covers six months prior to training start month and the six months post the training end month for each factory. Days with 0 efficiency are dropped from the analysis as these are reporting errors. Lines for which more than 20% of the days have zero efficiency for any of the three periods are dropped from analysis.

C.5.9 Treatment Effects on Supervisor Salary

To assess the effects of treatment on salary growth, we use the following specification:

$$\%growth_i = \alpha + \beta_1 T_i + \beta_2 NumMonths_i + \mu_s + \epsilon_i$$

where $\%growth_i$ is the percent change in gross salary between January 2017 and May 2018 (or the latest month observed) for the supervisor i, T_i is the treatment indicator, and $NumMonths_i$ is the number of months after January 2017 the supervisor is in the data (with a maximum of 18 if the supervisor is with the firm until May 2018), and μ_s is randomization strata fixed effects.

Results are reported in Appendix Table C.17. We find that treated supervisors experience 0.9 percentage points higher salary growth (7% on a baseline of 12.6 percentage points). In column 2, we additionally control for the number of months that elapses between January 2017 and the latest salary month available (number of months can at most be 18 if the supervisor is still with the firm until May 2018). As the change in salary would be increasing with time before quitting, we control for the number of months in order to control for the retention effects of training. Treatment increases the percent change in salary by 0.8 percentage points after controlling for retention effects (6 % of baseline). We do not find this effect to be heterogeneous by middle manager recommendation.

Table C.17: Treatment Effects on Salary Progression

	Salary Change	
	(1)	(2)
Treated	0.009**	0.008**
	(0.004)	(0.004)
Num. Months Before Quitting		0.014***
		(0.001)
Observations	1411	1411
Control Mean of Dependent Variable	.126	.126
Strata FE	X	X

Note: *** p<0.01, ** p<0.05, * p<0.10. The monthly salary data covers January 2017 to May 2018. For each supervisor, the percent change in salary is calculated as the percent change from the earliest to latest gross salary recorded. Supervisors who quit between January 2017 and training start are dropped from the analysis.

C.5.10 Treatment Effects on Incentive Bonuses

We aggregate the daily data on individual incentive payments to line-day level by summing up the individual payment amounts. We employ a specification parallel to our main difference-in-difference specification in 3.1, with incentive payments as the outcome variable. We have two outcomes of interest. First, we focus on the extensive margin of bonus payments by looking at an indicator for whether incentive payments have been made on the floor on a given day. Second, we use the inverse hyperbolic sine (IHS) transformation of the payment amount to explore the effects on magnitude of incentive payments.

Table C.18: Treatment Effects on Incentive Payments

	Sample: All Employees				Sample: Non - Supervisors			
	1[Any]		**IHS(Amount)**		**1[Any]**		**IHS(Amount)**	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
During X Treatment	0.034*	0.031*	0.269*	0.241	0.033*	0.031	0.263*	0.236
	(0.019)	(0.019)	(0.155)	(0.155)	(0.019)	(0.019)	(0.154)	(0.154)
After X Treatment	0.047*	0.041*	0.377*	0.333*	0.046*	0.041*	0.372*	0.328*
	(0.024)	(0.024)	(0.200)	(0.199)	(0.024)	(0.024)	(0.198)	(0.197)
Observations	270661	270661	270661	270661	270661	270661	270661	270661
Num. Lines	476	476	476	476	476	476	476	476
Cont. Mean	.081	.081	.65	.65	.081	.081	.646	.646
Line FE	X	X	X	X	X	X	X	X
Day FE	X	X	X	X	X	X	X	X
Relative Day FE		X		X		X		X

Note: *** p<0.01, ** p<0.05, * p<0.10. Standard errors are clustered at line level. $\mathbb{1}[Any]$ indicates any incentive payments have been paid in the line on a given day. $IHS(Amount)$ is the inverse hyperbolic sine transformation of the total incentive payments in the line on a given day.

Appendix Table C.18 shows the results, with columns 2 and 4 using our preferred specification. On the extensive margin, we find that the training increases the probability of having any bonus payments on the line by 3 percentage points during and 4 percentage points six months after the training (significant at 10%). These are large magnitudes as they represent a 38% and 51% increase from the control mean. On the intensive margin, we find that lines with all treated supervisors have 26% increase in incentive payments during the training period (not statistically significant), and 37% increase six months after training (significant at 10 %).[9]. In columns 5-8, we replicate the same analysis, but only focus on incentive payments to employees who are not supervisors or managers to assess the impact of training on workers.[10] The results are very similar to the results using the full sample, indicating that the effects also accrue to the workers, not just the supervisors who have been trained or to managers.

[9]Bellemare and Wichman (2019) notes that IHS-linear specifications with dummy variables can be interpreted similarly to log-linear specifications under conditions that our setting satisfies. Therefore, we calculate the approximate percentage change from treating all the supervisors on the line using the formula $e^{\hat{\beta}-0.5\hat{V}(\hat{\beta})}-1$ where $\hat{\beta}$ is the coefficient of interest and $\hat{V}()$ is the estimate of the variance.

[10]Specifically, we exclude any employee whose designation includes the words "supervisors", "manager", "senior executive", and "floor incharge".

C.5.11 Treatment Effects on Supervisor Attendance

Using the attendance roster, we assess the day-supervisor level retention effects of training using the following difference-in-differences specification:

$$\mathbb{1}[Attended]_{itr} = \alpha + \beta_1 T_i \, \mathbb{1}[During_t] + \beta_2 T_i \, \mathbb{1}[Post_t] + \delta_i + \mu_t + \gamma_r + \epsilon_{itr} \qquad \text{(C.8)}$$

where $\mathbb{1}[Attended]_{itr}$ is an indicator for whether supervisor i attended work on date t, T_i is the treatment indicator, and the $\mathbb{1}[During_t]$ and $\mathbb{1}[Post_t]$ are indicators for whether training is ongoing or over in the factory of the supervisor, and δ_i are supervisor fixed effects. The results are shown in Append Table C.19. We do not observe any evidence of treatment effects on supervisor retention.

Table C.19: Treatment Effects on Supervisor Attendance

	Daily Attendance		
	(1)	(2)	(3)
Treatment X During	0.001	0.001	0.001
	(0.005)	(0.005)	(0.005)
Treatment X Post	0.004	0.004	0.003
	(0.009)	(0.009)	(0.009)
Observations	516805	516805	516805
Number of Supervisors	1636	1636	1636
Control Mean of Dependent Variable	.895	.895	.895
Supervisor FE	X	X	X
Date FE		X	X
Relative Date FE			X

Note: *** p<0.01, ** p<0.05, * p<0.10. Outcome variable is the daily attendance of a supervisor. Standard errors are clustered at the supervisor level.

C.5.12 Treatment Effects on Worker Retention and Attendance

Appendix Figure C.10 shows survival curves for quitting for workers in lines with at least one supervisor treated versus none. There is no evidence of differential retention. Running a cox proportional hazard model with the preferred treatment definition of fraction of supervisors treated also yields no evidence of differential retention. For attendance, we follow an analogous approach to equation 7 for estimating the treatment effects on worker attendance, except with continuous line-level treatment. We do not find any evidence for treatment effects.

Figure C.10: Worker Retention by Supervisor Treatment

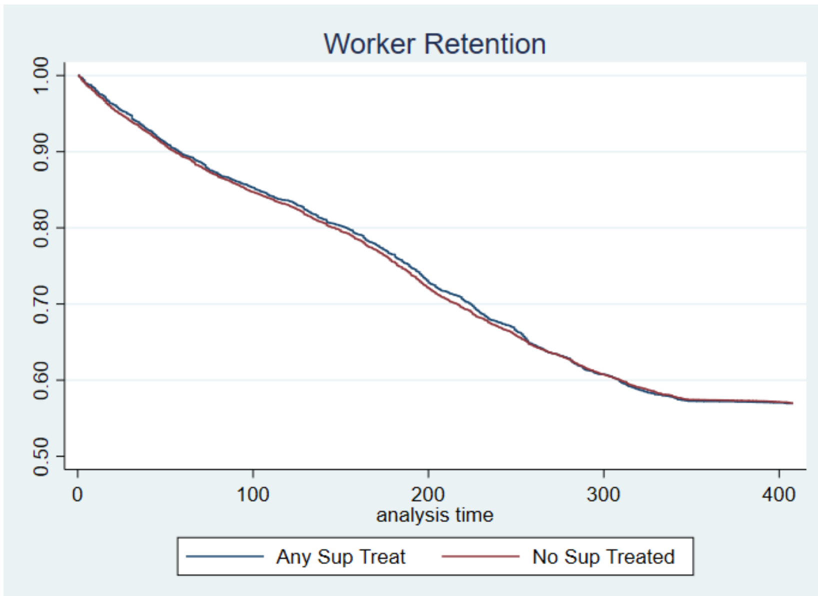

Table C.20: Treatment Effects on Worker Attendance

	Daily Attendance			
	(1)	(2)	(3)	(4)
During Training X Treatment	-0.001	-0.001	0.002	-0.000
	(0.004)	(0.004)	(0.004)	(0.002)
Post Training X Treatment	0.005	0.005	0.009	-0.001
	(0.006)	(0.006)	(0.005)	(0.004)
Observations	10864000	10864000	10864000	10863731
Cont. Mean of Dep. Var.	.86	.86	.86	.86
Line FE	X	X	X	
Employee FE				X
Day FE		X	X	X
Relative Day FE			X	X

Note: Standard errors are clustered at line level. *** p<0.01, ** p<0.05, * p<0.10. Coefficients are from a linear probability model on whether the employee has attended work on a given day. Sundays and days where less than 40% of employees attend work are dropped from the analysis.

C.6 Observable Determinants of Middle Manager Recommendation in Random Subsamples of Middle Managers

In this subsection, we rerun the LASSO analysis described in Section 3.5.1.2 and Appendix Section C.4.2 on random subsets of middle managers to assess any evidence of distinct and possibly countervailing ranking strategies by middle managers, especially with regards to their skill scores for the supervisors. To do this, we randomly select 50% of middle managers in our data, rerun the analysis with the same large set of possible determinants, and collect the selected variables. We repeat this exercise 1000 times.

With 50% subsamples, the LASSO analysis tends to select more variables and have higher R^2 than the analysis on the full sample. While analysis on the full sample selects 11 variables that, when included in a regression explaining the middle manager recommendation, give an R^2 of around 10%, the subsample analyses on average select 37 variables and produce an R^2 of 32% (both mean and median). However, we still interpret this relatively small R^2 on average as evidence that even in subsamples with potentially more homogeneous ranking strategies among middle managers we struggle to explain even half the variation in middle manager rankings with a very rich set of observables. Indeed the R^2 never reaches 50% in any of the 1000 iterations.

To identify observables that capture dimensions that middle managers might use to recommend supervisors in potentially countervailing ways, we look for observables that, when selected by the LASSO procedure, have both positive and negative coefficients a substantial fraction of the runs. Of course, if an observable has no relationship with the training recommendation, it would have positive and negative coefficients each about half the time due to noise. Therefore, we focus on observables that get selected by the LASSO at least half the time, and, when they are selected, they are significantly (at 5%) associated with the middle manager recommendation at least half the time.[11] The idea is that, if, for example, half the middle managers use an observable to positively recommend supervisors and the other half negatively, certain random samples that favor one group or the other due to chance would provide significant associations. In the case where there is no relationship between the characteristic and recommendation, this variable should not be significantly associated with the recommendation in the subsamples. This leaves us with 16 variables out of 50. All of these variables are either negative or positive over 90% of the runs in which they are

[11] For each run, we run a linear regression of middle manager recommendation on all the selected variables and use the statistical significance of the coefficient from this regression.

selected, providing no indication that there is substantial countervailing heterogeneity or a bimodal distribution of coefficient signs in how middle managers recommend based on these characteristics.

Further, the variables that are most often selected by the LASSO have similar interpretation to the results from the pooled sample. For example, middle manager perceptions of supervisors motivation, supervisor effort (as measured by target effort index), and the supervisors quantity and variety of experience (tenure at line, whether operator before, worked different line before) are consistently selected and often significant but essentially always positive such that to the degree that middle managers rank on this criteria they always reward motivation and tenure positively. Similarly, cognitive ability is often selected and consistently with a negative coefficient, consistent with the pooled analysis. Taken together we see no clear evidence of countervailing ranking strategies for any variables which are often selected and/or significant. Further, we also see that the set of variables which are often selected and significant is generally similar to those from the full sample exercise and still amount to a relatively small explanatory power (though larger than the pooled analysis) even in this analysis across smaller subsamples.

Specifically looking at skill scores given by middle managers, baseline management skills is most consistently selected by the LASSO analysis (though it ranks 20 out of 50 variables). This is followed by baseline industrial engineering skills (at rank 39). When selected, both these variables have consistently negative coefficients (around 90% of the time), suggesting that middle managers tend to allocate training to supervisors lacking these skills at baseline. That these variables are not selected in the pooled sample and are selected relatively less often in this subsample analysis likely reflects their low explanatory power as opposed to strong opposing ranking choices by middle managers. Finally, supervisor technical tailoring skill score is among the least selected variables, and, even when selected, is significant only a quarter of the time. As opposed to other skills, it has a positive coefficient 75 % of the times it is significant.

BIBLIOGRAPHY

Abarcar, P. (2019): "The Return Motivations of Legal Permanent Migrants: Evidence from Exchange Rate Shocks and Immigrants in Australia," *Journal of Economic Behavior and Organization*.

Abarcar, P. and C. Theoharides (2022): "Medical Worker Migration and Origin-Country Human Capital: Evidence from U.S. Visa Policy," *Review of Economics and Statistics*.

Acemoglu, D. (1997): "Training and innovation in an imperfect labour market," *The Review of Economic Studies*, 64, 445–464.

Acemoglu, D., P. Aghion, C. Lelarge, J. Van Reenen, and F. Zilibotti (2007): "Technology, Information, and the Decentralization of the Firm," *The Quarterly Journal of Economics*, 122, 1759–1799.

Acemoglu, D. and J.-S. Pischke (1998): "Why Do Firms Train? Theory and Evidence," *The Quarterly Journal of Economics*, 113, 79–119.

——— (1999): "Beyond Becker: Training in imperfect labour markets," *The Economic Journal*, 109, 112–142.

Acosta, P., C. Calderon, P. Fajnzylber, and H. Lopez (2008): "What is the Impact of International Remittances on Poverty and Inequality in Latin America," *World Development*, 36.

Adao, R., M. Kolesár, and E. Morales (2019a): "Shift-Share Designs: Theory and Inference," *The Quarterly Journal of Economics*, 134, 1949–2010.

Adao, R., M. Kolesar, and E. Morales (2019b): "Shift-Share Designs: Theory and Inference," *Quartery Journal of Economics*, 134.

Adhikari, S., S. Chaudhary, and N. L. Ekeator (2021): *Of Roads Less Traveled: Assessing the Potential of Economic Migration to Provide Overseas Jobs for Nigeria's Youth*, The World Bank, Washngton, DC.

Adhvaryu, A., V. Bassi, A. Nyshadham, J. Tamayo, and N. Torres (2023): "Organizational Responses to Product Cycles," *Available at SSRN 4403515*.

Adhvaryu, A., V. Bassi, A. Nyshadham, and J. A. Tamayo (2020): "No line left behind: Assortative matching inside the firm," Tech. rep., NBER.

Adhvaryu, A., N. Kala, and A. Nyshadham (2022a): "Management and Shocks to Worker Productivity," *Journal of Political Economy*, 130, 1–47.

——— (2022b): "Returns to On-the-job Soft Skills Training," *Journal of Political Economy*, forthcoming.

Adhvaryu, A., A. Nyshadham, and J. Tamayo (2022c): "An Anatomy of Performance Monitoring," .

——— (2022d): "Managerial Quality and Productivity Dynamics," *The Review of Economic Studies*, forthcoming.

Aghion, P., N. Bloom, B. Lucking, R. Sadun, and J. Van Reenen (2021): "Turbulence, firm decentralization, and growth in bad times," *American Economic Journal: Applied Economics*, 13, 133–69.

Aghion, P., N. Bloom, and J. Van Reenen (2014): "Incomplete contracts and the internal organization of firms," *The Journal of Law, Economics, & Organization*, 30, i37–i63.

Agte, P., A. Bernhardt, E. Field, R. Pande, and N. Rigol (2022): "Investing in the Next Generation: The Long-Run Impacts of a Liquidity Shock," *NBER Working Paper*.

Agunias, D. R. (2010): "Migration's Middlemen: Regulating Recruitment Agencies in the Philippines-United Arab Emirates Corridor," https://www.migrationpolicy.org/research/migrations-middlemen-regulating-recruitment-agencies-philippines-united-arab-emirates.

Akram, A., S. Chowdhury, and A. Mobarak (2017): "Effects of Emigration on Rural Labor Markets," *Working Paper*.

Alan, S., G. Corekcioglu, and M. Sutter (2023): "Improving workplace climate in large corporations: A clustered randomized intervention," *The Quarterly Journal of Economics*, 138, 151–203.

Alfonsi, L., O. Bandiera, V. Bassi, R. Burgess, I. Rasul, M. Sulaiman, and A. Vitali (2020): "Tackling youth unemployment: Evidence from a labor market experiment in Uganda," *Econometrica*, 88, 2369–2414.

Allen, T., C. Arkolakis, and Y. Takahashi (2020): "Universal Gravity," *Journal of Political Economy*, 128, 393–433.

Asis, M. and D. R. Aguinas (2012): "Strengthening Pre-Departure Orientation Programmes in Indonesia, Nepal, and the Philippines," *Migration Policy Institute Issue in Brief No. 5.*

Atkin, D. (2016): "Endogenous skill acquisition and export manufacturing in Mexico," *American Economic Review*, 106, 2046–2085.

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017): "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan*," *The Quarterly Journal of Economics*, 132, 1101–1164.

Balboni, C., O. Bandiera, R. Burgess, M. Ghatak, and A. Heil (2021): "Why Do People Stay Poor?" *Quarterly Journal of Economics*.

Bandiera, O., M. C. Best, A. Q. Khan, and A. Prat (2021): "The allocation of authority in organizations: A field experiment with bureaucrats," *The Quarterly Journal of Economics*, 136, 2195–2242.

Bandiera, O., R. Burgess, N. Das, S. Gulesci, I. Rasul, and M. Sulaiman (2017): "Labor Markets and Poverty in Village Economies," *Quarterly Journal of Economics*.

Bandiera, O., A. Prat, S. Hansen, and R. Sadun (2020): "CEO behavior and firm performance," *Journal of Political Economy*, 128, 1325–1369.

Banerjee, A. and E. Duflo (2014): "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program," *Review of Economic Studies*, 81, 572–607.

Banerjee, A., E. Duflo, N. Goldberg, D. Karlan, R. Osei, W. Pariente, J. Shapiro, B. Thuysbaert, and C. Udry (2015): "A multifaceted program causes lasting progress for the very poor: Evidence from six countries," *Science*, 348.

Banerjee, A., E. Duflo, and G. Sharma (2021): "Long-Term Effects of the Targeting the Ultra Poor Program," *American Economic Review: Insights*, 4, 471–486.

Barrera-Osorio, F., A. D. Kugler, and M. I. Silliman (2020): "Hard and soft skills in vocational training: Experimental evidence from Colombia," Tech. rep., NBER.

Barsbai, T., A. Steinmayr, and C. Winter (2019): "Immigration into a Recession," Tech. rep., mimeo.

Batista, C., A. Lacuesta, and P. Vicente (2012): "Testing the brain gain hypothesis: micro evidence from Cape Verde," *Journal of Development Economics*.

Bazzi, S. (2017): "Wealth Heterogeneity and the Income Elasticity of Migration," *American Economic Journal: Applied Economics*, 9, 219–255.

Bazzi, S., L. Cameron, S. Schaner, and F. Witoelar (2021a): "Information, Intermediaries, and International Migration," *NBER Working Paper No. 29588*.

Bazzi, S., L. Cameron, S. G. Schaner, and F. Witoelar (2021b): "Information, Intermediaries, and International Migration," .

Beam, E. A., D. McKenzie, and D. Yang (2016): "Unilateral Facilitation Does Not Raise International Labor Migration from the Philippines," *Economic Development and Cultural Change*, 64, 323–368.

Becker, G. S. (1964): *Human Capital*, New York: Columbia University Press.

Beine, M., F. Docquier, and Çağlar Özden (2011): "Diasporas," *Journal of Development Economics*, 95, 30–41, symposium on Globalization and Brain Drain.

Beine, M. and C. R. Parsons (2017): "Climatic Factors as Determinants of International Migration: Redux," *CESifo Economic Studies*, 63, 386–402.

Bellemare, M. and C. Wichman (2019): "Elasticities and the Inverse Hyperbolic Sine Transformation," *Working Paper*.

Benson, A., D. Li, and K. Shue (2019): "Promotions and the Peter Principle*," *The Quarterly Journal of Economics*, 134, 2085–2134.

Berlemann, M. and M. F. Steinhardt (2017): "Climate Change, Natural Disasters, and Migration—a Survey of the Empirical Evidence," *CESifo Economic Studies*, 63, 353–385.

Bertoli, S. and J. Fernández-Huertas Moraga (2013): "Multilateral Resistance to Migration," *Journal of Development Economics*, 102, 79–100.

Bertoli, S., J. Fernández-Huertas Moraga, and S. Keita (2017): "The Elasticity of the Migrant Labour Supply: Evidence from Temporary Filipino Migrants," *The Journal of Development Studies*, 53, 1822–1834.

Bertrand, M. and A. Schoar (2003): "Managing with style: The effect of managers on firm policies," *The Quarterly journal of economics*, 118, 1169–1208.

Bianchi, N. and M. Giorcelli (2022): "The dynamics and spillovers of management interventions: Evidence from the training within industry program," *Journal of Political Economy*, 130, 1630–1675.

Bilinski, A., J. Poe, J. Roth, and P. Sant'Anna (2022): "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," Tech. rep.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2013): "DOES MANAGEMENT MATTER? EVIDENCE FROM INDIA," *The Quarterly Journal of Economics*, 1, 51.

Bloom, N., R. Lemos, R. Sadun, D. Scur, and J. Van Reenen (2014): "JEEA-FBBVA Lecture 2013: The new empirical economics of management," *Journal of the European Economic Association*, 12, 835–876.

Bloom, N., R. Sadun, and J. Van Reenen (2016): "Management as a Technology?" Tech. rep., NBER.

Bloom, N. and J. Van Reenen (2007): "Measuring and Explaining Management Practices across Firms and Countries," *The Quarterly Journal of Economics*, 1351–1408.

——— (2011): "Human resource management and productivity," *Handbook of labor economics*, 4, 1697–1767.

Blumenstock, J. E., N. Eagle, and M. Fafchamps (2016): "Airtime Transfers and Mobile Communications: Evidence in the Aftermath of Natural Disasters," *Journal of Development Economics*, 120, 157–181.

Borjas, G. J. (2003): "The Labor Demand Curve Is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market," *The Quarterly Journal of Economics*, 118, 1335–1374.

Borusyak, K. and P. Hull (2020): "Non-Random Exposure to Exogenous Shocks: Theory and Applications," Tech. rep., National Bureau of Economic Research.

Borusyak, K., P. Hull, and X. Jaravel (2022a): "Quasi-Experimental Shift-Share Designs," *Review of Economic Studies*, 89, 181–213.

——— (2022b): "Quasi-Experimental Shift-Share Research Designs," *The Review of Economic Studies*, 89, 181–213.

Borusyak, K., X. Jaravel, and J. Spiess (2022c): "Revisiting Event Study Designs: Robust and Efficient Estimation," Tech. rep.

Bossavie, L., J.-S. Gorlach, C. Ozden, and H. Wang (2021a): *Temporary Migration for Long-Term Investment*, Policy Research Working Papers, The World Bank.

——— (2021b): "Temporary Migration for Long-Term Investment," *World Bank Policy Research Working Paper*.

Bossavie, L. and Ç. Özden (2023): "Impacts of Temporary Migration on Development in Origin Countries," *The World Bank Research Observer*, lkad003.

Breza, E. and C. Kinnan (2021): "Measuring the Equilibrium Impacts of Credit: Evidence from the Indian Microfinance Crisis," *The Quarterly Journal of Economics*, 136, 1447–1497.

Bryan, G., S. Chowdhury, and A. M. Mobarak (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82, 1671–1748.

Bryan, G. and M. Morten (2019): "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*.

Caballero, M. E., B. Cadena, and B. K. Kovak (2021): "The International Transmission of Local Economic Shocks Through Migrant Networks," .

Caballero, M. E., B. C. Cadena, and B. K. Kovak (2023): "The International Transmission of Local Economic Shocks Through Migrant Networks," *Journal of International Economics*, 145.

Cai, J. and A. Szeidl (2022): "Indirect Effects of Access to Finance," *NBER Working Paper*.

Caicedo, S., M. Espinosa, and A. Seibold (2022): "Unwilling to Train?—Firm Responses to the Colombian Apprenticeship Regulation," *Econometrica*, 90, 507–550.

Caliendo, L., G. Mion, L. D. Opromolla, and E. Rossi-Hansberg (2020): "Productivity and organization in Portuguese firms," *Journal of Political Economy*, 128, 4211–4257.

Caliendo, L., F. Monte, and E. Rossi-Hansberg (2015): "The anatomy of French production hierarchies," *Journal of Political Economy*, 123, 809–852.

Caliendo, L. and E. Rossi-Hansberg (2012): "The impact of trade on organization and productivity," *The quarterly journal of economics*, 127, 1393–1467.

Callaway, B., A. Goodman-Bacon, and P. H. Sant'Anna (2021): "Difference-in-Differences with a Continuous Treatment," *arXiv preprint arXiv:2107.02637*.

Callaway, B. and P. H. Sant'Anna (2021): "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 225, 200–230.

Cattaneo, C., M. Beine, C. J. Fröhlich, D. Kniveton, I. Martinez-Zarzoso, M. Mastrorillo, K. Millock, E. Piguet, and B. Schraven (2019): "Human Migration in the Era of Climate Change," *Review of Environmental Economics and Policy*, 13, 189–206.

Cattaneo, C. and G. Peri (2016): "The Migration Response to Increasing Temperatures," *Journal of Development Economics*, 122, 127–146.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019): "The Effect of Minimum Wages on Low-Wage Jobs*," *The Quarterly Journal of Economics*, 134, 1405–1454.

Chand, S. and M. A. Clemens (2019): "Human Captial Investment under Exit Options: Evidence from a Natural Quasi-Experiment," *IZA Discussion Papers No. 12173*.

Chioda, L., D. Contreras-Loya, P. Gertler, and D. Carney (2021): "Making entrepreneurs: Returns to training youth in hard versus soft business skills," Tech. rep., NBER.

Chiswick, B. and T. Hatton (2003): "International Migration and the Integration of Labor Markets," NBER Chapters, National Bureau of Economic Research, Inc.

Choi, H. and D. Yang (2007): "Are Remittances Insurance? Evidence from Rainfall Shocks in the Philippines," *The World Bank Economic Review*, 21, 219–248.

Cinque, A. and L. Reiners (2023): "Confined to Stay: Natural Disasters and Indonesia's Migration Ban," .

Clemens, M. (2010): *The Annual Proceedings of the Wealth and Well-Being of Nations: 2009-2010*, Beloit College Press, Beloit, WI, chap. The Biggest Idea in Development That No One Really Tried, 25–50.

Clemens, M., C. Montenegro, and L. Pritchett (2019): "The Place Premium: Bounding the Price Equivalent of Migration Barriers," *Review of Economics and Statistics*.

Clemens, M. and L. Pritchett (2013): "Time-Bound Labor Access to the United States: A Four-Way Win for the Middle Class, Low-Skill Workers, Border Security, and Migrants," *Center for Global Development Brief*.

Clemens, M. A. (2011a): "Economics and Emigration: Trillion-Dollar Bills on the Sidewalk?" *Journal of Economic Perspectives*, 25, 83–106.

——— (2011b): "Economics and Emigration: Trillion-Dollar Bills on the Sidewalk?" *The Journal of Economic Perspectives*, 25, pp. 83–106.

Clemens, M. A. and E. R. Tiongson (2017): "Split Decisions: Household Finance When a Policy Discontinuity Allocates Overseas Work," *The Review of Economics and Statistics*, 99, 531–543.

Conley, T. G. (1999): "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 92, 1–45.

Convergences (2019): *Microfinance Barometer 2019: A Look Back at the Trends in Microfinance*, Convergences.

Cortes, P. (2015): "The Feminization of International Migration and Its Effects on the Children Left Behind: Evidence from the Philippines," *World Development*, 65, 62–78.

CRED (2017): "Annual Disaster Statistical Review 2017," .

Cuadros-Menaca, A. and A. Gaduh (2020): "Remittances, Child Labor, and Schooling: Evidence from Colombia," *Economic Development and Cultural Change*, 68, 1257–1293.

Dal Bó, E., F. Finan, N. Y. Li, and L. Schechter (2021): "Information Technology and Government Decentralization: Experimental Evidence from Paraguay," *Econometrica*, 89, 677–701.

De Arcangelis, G., A. Fertig, Y. Liang, P. Srouji, and D. Yang (2023): "Measuring Remittances," *Journal of Development Economics*, 161, 103004.

de Brauw, A. and J. Giles (2017): "Migrant Opportunity and the Educational Attainment of Youth in Rural China," *Journal of Human Resources*, 52.

de Chaisemartin, C. and X. D'Haultfœuille (2020): "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 110, 2964–96.

de Mel, S., D. McKenzie, and C. Woodruff (2008): "Returns to Capital in Microenterprises: Evidence from a Field Experiment," *The Quarterly Journal of Economics*, 123, 1329–1372.

Deming, D. J. (2022): "Four facts about human capital," *Journal of Economic Perspectives*, 36, 75–102.

Deryugina, T. (2017): "The Fiscal Cost of Hurricanes: Disaster Aid versus Social Insurance," *American Economic Journal: Economic Policy*, 9, 168–98.

Deserranno, E., S. A. Caria, G. Leon-Ciliotta, and P. Kastrau (2022): *The allocation of incentives in multi-layered organizations*, Universitat Pompeu Fabra, Department of Economics and Business.

Dimitriadis, S. and R. Koning (2022): "Social Skills Improve Business Performance: Evidence from a Randomized Control Trial with Entrepreneurs in Togo," *Management Science*, 68, 8635–8657.

Dinkelman, T., G. Kumchulesi, and M. Mariotti (2024): "Labor Migration, Capital Accumulation, and the Structure of Rural Labor Markets," *Working Paper*.

——— (forthcoming): "Labor Migration, Capital Accumulation, and the Structure of Rural Labor Markets," *Review of Economics and Statistics*.

Dinkelman, T. and M. Mariotti (2016): "The Long Run Effect of Labor Migration on Human Capital Formation in Communities of Origin," *American Economic Journal: Applied Economics*.

Docquier, F. and H. Rapoport (2012): "Globalization, brain drain, and development," *Journal of Economic Literature*.

Drabo, A. and L. M. Mbaye (2015): "Natural Disasters, Migration and Education: An Empirical Analysis in Developing Countries," *Environment and Development Economics*, 20, 767–796.

Ducanes, G. (2010): "The Case of Missing Remittances in the FIES: Could it be Causing Us to Mismeasure Welfare Changes?" *University of the Philippines School of Economics Discussion Paper*.

Dustmann, C. and O. Kirchkamp (2002): "The optimal migration duration and activity choice after re-migration," *Journal of Development Economics*.

Dustmann, C., H. Ku, and T. Surovtseva (2021): *Real Exchange Rates and the Earnings of Immigrants*, Centre for Research and Analysis of Migration, Department of Economics ....

——— (2023): "Real Exchange Rates and the Earnings of Immigrants," *The Economic Journal*, 134, 271–294.

Dustmann, C. and J. Mestres (2010): "Remittances and Temporary Migration," *Journal of Development Economics*, 92, 62–70.

Dustmann, C., U. Schönberg, and J. Stuhler (2016): "The Impact of Immigration: Why Do Studies Reach Such Different Results?" *Journal of Economic Perspectives*, 30, 31–56.

Eaton, J. and S. Kortum (2002): "Technology, geography and trade," *Econometrica*, 70, 1741–1779.

Eckstein, Z. and Y. Weiss (2004): "On the Wage Growth of Immigrants: Israel, 1990–2000," *Journal of the European Economic Association*, 2, 665–695.

Edmonds, E. and C. Theoharides (2020): "The Short Term Impact of a Productive Asset Transfer in Families with Child Labor: Experimental Evidence from the Philippines." *Journal of Development Economics*, 146.

Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. W. Walker (2022): "General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya," *Econometrica*.

Emanuel, K. (2011): "Global Warming Effects on U.S. Hurricane Damage," *Weather, Climate, and Society*, 3, 261–268.

Fairburn, J. A. and J. M. Malcomson (2001): "Performance, Promotion, and the Peter Principle," *The Review of Economic Studies*, 68, 45–66.

Fernando, A. N. and N. Singh (2021): "Regulation by Reputation? Quality Revelation of Labor Intermediaries in International Migration," *Working Paper*.

——— (2023): "Regulation by Reputation? Intermediaries, Labor Abuses, and International Migration," .

Franklin, S. and J. Labonne (2019): "Economic Shocks and Labor Market Flexibility," *Journal of Human Resources*, 54, 171–199.

Frederiksen, A., L. B. Kahn, and F. Lange (2020): "Supervisors and performance management systems," *Journal of Political Economy*, 128, 2123–2187.

Gaikwad, N., K. Hanson, and A. Toth (2024): "Bridging the Gulf: How Migration Fosters Tolerance, Cosmopolitanism, and Support for Globalization," *Working Paper*.

Ghatak, M. (2015): "Theories of Poverty Traps and Anti-Poverty Policies," *World Bank Economic Review*, 29, S77–S105.

Giannelli, G. C. and E. Canessa (2022): "After the Flood: Migration and Remittances as Coping Strategies of Rural Bangladeshi Households," *Economic Development and Cultural Change*, 70, 1159–1195.

Gibson, J., D. McKenzie, H. Rohorua, and S. Stillman (2018): "The Long-term Impacts of International Migration: Evidence from a Lottery," *World Bank Economic Review*, 32, 127–47.

Gibson, J., D. McKenzie, and S. Stillman (2010): "How Important Is Selection? Experimental vs. Non-Experimental Measures of the Income Gains from Migration," *Journal of the European Economic Association*.

——— (2011): "The Impacts of International Migration on Remaining Household Members: Omnibus Result from a Migration Lottery Program," *Review of Economics and Statistics*, 93.

——— (2014): "The Development Impact of a Best Practice Seasonal Migration Policy," *Review of Economics and Statistics*.

Giorcelli, M. (2019): "The long-term effects of management and technology transfers," *American Economic Review*, 109, 121–52.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020): "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 110, 2586–2624.

Goodman-Bacon, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277, themed Issue: Treatment Effect 1.

Gosnell, G. K., J. A. List, and R. D. Metcalfe (2020): "The impact of management practices on employee productivity: A field experiment with airline captains," *Journal of Political Economy*, 128, 1195–1233.

Gröger, A. (2019): "Easy Come, Easy Go? Economic Shocks, Labor Migration and the Family Left Behind," *Journal of International Economics*, 128.

——— (2021): "Easy Come, Easy Go? Economic Shocks, Labor Migration and the Family Left Behind," *Journal of International Economics*, 128.

Groh, M., N. Krishnan, D. McKenzie, and T. Vishwanath (2016): "The impact of soft skills training on female youth employment: evidence from a randomized experiment in Jordan," *IZA Journal of Labor & Development*, 5, 1–23.

Groh, M., N. Krishnan, D. J. McKenzie, and T. Vishwanath (2012): "Soft skills or hard cash? The impact of training and wage subsidy programs on female youth employment in Jordan," *World Bank Policy Research Working Paper*.

Gröschl, J. and T. Steinwachs (2017): "Do Natural Hazards Cause International Migration?" *CESifo Economic Studies*, 63, 445–480.

Haegele, I. (2022a): "Talent Hoarding in Organization," *Available at SSRN 3977728*.

——— (2022b): "Talent Hoarding in Organizations," Tech. rep.

Halliday, T. (2006): "Migration, Risk, and Liquidity Constraints in El Salvador," *Economic development and cultural change*, 54, 893–925.

Hanson, G. and C. McIntosh (2012): "Birth Rates and Border Crossings: Latin American Migration to the US, Canada, Spain and the UK," *Economic Journal*, 122, 707–726.

Harris, R. D. and E. Tzavalis (1999): "Inference for unit roots in dynamic panels where the time dimension is fixed," *Journal of Econometrics*, 91, 201–226.

Hjort, J., H. Malmberg, and T. Schoellman (2022): "The missing middle managers: Labor costs, firm structure, and development," Tech. rep., NBER.

Hoffman, M. and S. Tadelis (2021): "People management skills, employee attrition, and manager rewards: An empirical analysis," *Journal of Political Economy*, 129, 243–285.

Hsiang, S. M. and A. S. Jina (2014): "The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones," .

Hussam, R., N. Rigol, and B. Roth (2022): "Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in the Field," *American Economic Review*, 112, 861–898.

IOM (2013): "Country Migration Report: The Philippines 2013," .

IPCC (2021): "Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change," .

Kaboski, J., M. Lipscomb, V. Midrigan, and C. Pelnik (2022): "How Important are Investment Indivisibilities for Investment? Empirical Evidence from Uganda," *NBER Working Paper*.

Karlan, D. and J. Morduch (2010): "Access to Finance," *Handbook of Development Economics*, 5, 4703–4784.

Karlan, D. and J. Zinman (2018): "Price and Control Elasticities of Demand for Savings," *Journal of Development Economics*, 130.

Khanna, G. and N. Morales (2023): "The IT boom and other unintended consequences of chasing the American dream," Working Paper.

Khanna, G., E. Murathanoglu, C. B. Theoharides, and D. Yang (2022): "Abundance from Abroad: Migrant Income and Long-Run Economic Development," Tech. rep., National Bureau of Economic Research.

Kinnan, C., S.-Y. Wang, and Y. Wang (2019): "Access to Migration for Rural Households," *American Economic Journal: Applied Economics*.

Kirdar, M. (2009): "Labor Market Outcomes, Savings Accumulation, and Return Migration," *Labour Economics*.

Kleemans, M. and J. Magruder (2019): "Labor Market Changes in Response to Immigration: Evidence from Internal Migration Driven by Weather Shocks," *The Economic Journal*.

Kossin, J. P., T. Hall, T. Knutson, K. E. Kunkel, R. J. Trapp, D. E. Waliser, and M. F. Wehner (2017): *Extreme storms*, Washington, D.C.: U.S. Global Change Research Program, 256–276.

Lagakos, D., A. M. Mobarak, and M. E. Waugh (2023): "The Welfare Effects of Encouraging Rural–Urban Migration," *Econometrica*, 91, 803–837.

Lazear, E. (2004): "The Peter Principle: A Theory of Decline," *Journal of Political Economy*, 112, S141–S163.

Lazear, E. P., K. L. Shaw, and C. T. Stanton (2015): "The value of bosses," *Journal of Labor Economics*, 33, 823–861.

Llull, J. (2018): "Immigration, Wages, and Education: A Labor Market Equilibrium Structural Model," *The Review of Economic Studies*, 85, 1852–1896.

Lopez-Cordoba, E. (2005): "Globalization, Migration, and Development: The Role of Mexican Migrant Remittances," *Economia*, 6.

Macchiavello, R., A. Menzel, A. Rabbani, and C. Woodruff (2020): "Challenges of change: An experiment promoting women to managerial roles in the bangladeshi garment sector," Tech. rep., NBER.

Mahajan, P. and D. Yang (2020a): "Taken by Storm: Hurricanes, Migrant Networks, and US Immigration," *American Economic Journal: Applied Economics*, 12, 250–277.

——— (2020b): "Taken by Storm: Hurricanes, Migrant Networks, and U.S. Immigration," *American Economic Journal: Applied Economics*.

Marchiori, L., J.-F. Maystadt, and I. Schumacher (2012): "The Impact of Weather Anomalies on Migration in Sub-Saharan Africa," *Journal of Environmental Economics and Management*, 63, 355–374.

Mbaye, L. M. and A. Drabo (2017): "Natural Disasters and Poverty Reduction: Do Remittances Matter?" *CESifo Economic Studies*, 63, 481–499.

Mckenzie, D. and H. Rapoport (2007): "Network Effects and the Dynamics of Migration and Inequality: Theory and Evidence from Mexico," *Journal of Development Economics*, 84, 1–24.

McKenzie, D. and H. Rapoport (2011): "Can migration reduce educational attainment? Evidence from Mexico," *Journal of Population Economics*, 24, 1331–1358.

McKenzie, D., C. Theoharides, and D. Yang (2014): "Distortions in the International Migrant Labor Market: Evidence from Filipino Migration and Wage Responses to Destination Country Economic Shocks," *American Economic Journal: Applied Economics*, 6, 49–75.

McKenzie, D. and C. Woodruff (2017): "Business practices in small firms in developing countries," *Management Science*, 63, 2967–2981.

Mendola, M. (2012): "Rural Out-migration and economic development at origin: A review of the evidence," *Journal of International Development*, 24, 102–122.

Mincer, J. (1962): "On-the-job training: Costs, returns, and some implications," *Journal of political Economy*, 70, 50–79.

Mishra, P. (2007): "Emigration and Wages in Source Countries: Evidence from Mexico," *Journal of Development Economics*, 82, 180–199.

Mishra, U. S. and S. I. Rajan (2010): "Managing Migration from India: Lessons from the Philippines," in *India Migration Report 2010*, Routledge India.

Mobarak, A. M., I. Sharif, and M. Shrestha (2023a): "Returns to International Migration: Evidence from a Bangladesh-Malaysia Visa Lottery," *American Economic Journal: Applied Economics*, 15, 353–88.

——— (2023b): "Returns to International Migration: Evidence from a Bangladesh Malaysia Visa Lottery," *American Economic Journal: Applied Economics*, 15, 353–388.

Morris, C. N. (1983): "Parametric Empirical Bayes Inference: Theory and Applications," *Journal of the American Statistical Association*, 78, 47–55.

Mountford, A. (1997): "Can a brain drain be good for growth in the source country?" *Journal of Development Economics*.

Munshi, K. (2003a): "Networks in the Modern Economy: Mexican Migrants in the U. S. Labor Market*," *The Quarterly Journal of Economics*, 118, 549–599.

——— (2003b): "Networks in the Modern Economy: Mexican Migrants in the U. S. Labor Market," *The Quarterly Journal of Economics*, 118, pp. 549–599.

Naidu, S., Y. Nyarko, and S.-Y. Wang (2016): "Monopsony Power in Migrant Labor Markets: Evidence from the United Arab Emirates," *Journal of Political Economy*, 124, 1735–1792.

——— (2023): "The Benefits and Costs of Guest Worker Programs: Experimental Evidence from the India-UAE Migration Corridor," .

Nekoei, A. (2013): "Immigrants' Labor Supply and Exchange Rate Volatility," *American Economic Journal: Applied Economics*, 5.

Nunn, N. (2019): "Rethinking Economic Development," *Canadian Journal of Economics*, 54, 1349–1373.

Orefice, G., H. Rappaport, and G. Santoni (2023): "How Do Immigrants Promote Exports? Networks, Knowledge, Diversity," .

Orrenius, P., M. Zavodny, J. Canas, and R. Coronado (2010): "Do Remittances Boost Economic Development? Evidence from Mexican States," *Law and Business Review of the Americas*, 16, 803–822.

Park, Y. C. and J. W. Lee (2002): "Financial Crisis and Recovery: Patterns of Adjustment in East Asia, 1996-99," *Asian Development Bank*.

Pesaran, M. H. (2006): "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure," *Econometrica*, 74, 967–1012.

Porcher, C. (2022): "Migration with Costly Information," .

Pritchett, L. and F. Hani (2020): "The Economics of International Wage Differentials and Migration," *Oxford Research Encyclopedia of Economics and Finance*.

Radelet, S. and J. Sachs (2000): *The Onset of the East Asian Financial Crisis*, University of Chicago Press, 105–153.

Rigol, N. and B. N. Roth (2021): "Loan Officers Impede Graduation from Microfinance: Strategic Disclosure in a Large Microfinance Institution," Tech. rep., NBER.

Rosenzweig, M. R. and C. Udry (2019): "External Validity in a Stochastic World: Evidence from Low-Income Countries," *The Review of Economic Studies*, 87, 343–381.

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023): "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*.

Samson, Feng Chi Hsu, G. C. E. L. N. Z. (2021): "Improving Accuracy of Economic Estimations with VIIRS DNB Image Products," in *Remote Sensing of Night-time Light*, Routledge.

Sandvik, J., R. Saouma, N. Seegert, and C. T. Stanton (2021): "Should Workplace Programs be Voluntary or Mandatory? Evidence from a Field Experiment on Mentorship," Tech. rep., NBER.

Shrestha, M. (2019): "Death Scares: How Potential Work-Migrants Infer Mortality Rates from Migrant Deaths," *Journal of Development Economics*, 141, 102368.

——— (2020): "Get Rich or Die Tryin': Perceived Earnings, Perceived Mortality Rates, and Migration Decisions of Potential Work Migrants from Nepal," *World Bank Economic Review*, 34.

Shrestha, S. (2017): "No Man Left Behind: Effects of Emigration Prospects on Educational and Labor Outcomes of Non-migrants," *Economic Journal*.

Shrestha, S. and D. Yang (2019a): "Facilitating Worker Mobility: A Randomized Information Intervention among Migrant Workers in Singapore," *Economic Development and Cultural Change*, 68, 63–91.

Shrestha, S. A. and D. Yang (2019b): "Facilitating Worker Mobility: A Randomized Information Intervention among Migrant Workers in Singapore," *Economic Development and Cultural Change*, 68, 63–91.

Skoufias, E., E. Strobl, and T. Tveit (2021): "Can We Rely on VIIRS Nightlights to Estimate the Short-Term Impacts of Natural Hazards? Evidence from Five South East Asian Countries," *Geomatics, Natural Hazards and Risk*, 12, 381–404.

Stark, O., C. Helmenstein, and A. Prskawetz (1997): "A Brain Gain with a Brain Drain," *Economic Letters*.

Stogdill, R. M. and A. E. Coons (1957): "Leader behavior: Its description and measurement." .

Strobl, E. (2019): "The Impact of Typhoons on Economic Activity in the Philippines: Evidence from Nightlight Intensity," .

Sun, L. and S. Abraham (2021a): "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 225, 175–199.

——— (2021b): "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 225, 175–199.

Tang, S. H. K., Y. Wang, and Y. Wang (2022): "Curse of Lower-skilled Emigration on Human Capital Formation: Evidence from the Migration Surge of the 2000s," *Working Paper*.

Theoharides, C. (2018a): "Manila to Malaysia, Quezon to Qatar International Migration and Its Effects on Origin-Country Human Capital," *Journal of Human Resources*, 53, 1022–1049.

——— (2018b): "Manila to Malaysia, Quezon to Qatar: International Migration and the Effects on Origin-Country Human Capital," *Journal of Human Resources*.

——— (2020): "The Unintended Consequences of Migration Policy on Origin-Country Labor Market Decisions," *Journal of Development Economics*.

Tombe, T. and X. Zhu (2019): "Trade, Migration and Productivity: A Quantitative Analysis of China," *American Economic Review*, 109, 1843–1872.

United Nations (2019): "World Population Policies 2019," .

United Nations, T. (2019a): "International Migrant Stock 2019," Tech. rep., Department of Economic and Social Affairs, United Nations.

——— (2019b): "World Population Policies 2019: International Migration Policies and Programs," Tech. rep., United Nations.

Uprety, S., C. Cao, Y. Gu, X. Shao, S. Blonski, and B. Zhang (2019): "Calibration Improvements in S-NPP VIIRS DNB Sensor Data Record Using Version 2 Reprocessing," *IEEE Transactions on Geoscience and Remote Sensing*, 57, 9602–9611.

Winter, C. (2020): "Natural Disasters, Networks, and Migration," .

World Bank, T. (2018a): *Moving for Prosperity: Global Migration and Labor Markets*, Washington, DC: World Bank Publications.

——— (2018b): *State of the Social Safety Nets 2018*, World Bank.

——— (2023): "Leveraging Diaspora Finances for Private Capital Mobilization," *Migration and Development Brief*, 39.

Yang, D. (2006): "Why do migrants return to poor countries? Evidence from Philippine migrants' responses to exchange rate shocks," *The Review of Economics and Statistics*, 88, 715–735.

——— (2008a): "Coping with Disaster: The Impact of Hurricanes on International Financial Flows, 1970-2002," *The BE Journal of Economic Analysis & Policy*, 8.

——— (2008b): "Coping with Disaster: The Impact of Hurricanes on International Financial Flows, 1970-2002," *B.E. Journal of Economic Analysis and Policy (Advances)*, 8.

——— (2008c): "International Migration, Remittances, and Household Investment: Evidence from Philippine Migrants' Exchange Rate Shocks," *Economic Journal*, 118, 591–630.

——— (2008d): "Risk, Migration, and Rural Financial Markets: Evidence from Earthquakes in El Salvador," *Social Research: An International Quarterly*, 75, 955–992.

——— (2011): "Migrant Remittances," *Journal of Economic Perspectives*, 25, 129–52.