Three Essays in International Trade and Macroeconomics

by

Lin Ma

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2014

Doctoral Committee:

Associate Professor Andrei A. Levchenko, Chair
Professor Rüdiger Bachmann
Professor Alan V. Deardorff
Assistant Professor Kyle Handley

To my wife, Zheng Mu

ACKNOWLEDGEMENTS

I first want to thank my advisors Andrei Levchenko, Rüdiger Bachmann, Alan Deardorff, and Kyle Handley for their endless support, advice, and encouragement. They have guided me into the world of economic research, and they have never ceased to push me into a better economist. They have spent countless hours with me discussing ideas, reading and revising drafts, and challenging me to think harder. These interactions sharpened my analytical skills, broadened my view of the world, and eventually, made me a better researcher. I have greatly enjoyed my intellectual life in Ann Arbor, and learned tremendously at the same time — and none of these would have been possible without them.

I also want to thank my friends in Ann Arbor for their support. Fan Fei has given me many useful feed-backs on my papers, and helped me greatly during my job search. Many of my friends, such as Qing Lai, Xiang Zhou, Jin Li, Yiwen Zhou, and Siwei Cheng, study in the fields of sociology, anthropology, and linguistics. Talking with scholars from these exciting and colorful fields of social sciences has made the study of our own dismal science a lot less bleak than it is supposed to be, and to this end, I am grateful to their friendship. Special thank goes to my old roommate and friend, Yilin Mo, who has introduced me into many geeky corners of computer science, which eventually made my own coding days a lot easier.

Finally I want to thank my family. My parents are my first teachers. They have shaped me into what I am today, and are still making me into a better person with their love, support, and sacrifice. The last thank goes to my wife, Zheng, who has endured all the hardships and shared all the laughters with me all these years. She always has faith in me, and it is her encouragement and companionship that guided me through the difficult times. She is the master of language between the two of us, and she has proofread many of my papers, usually more than once. I am deeply grateful for this. However, I also want to shamelessly point out that, the remaining typos you will find in this dissertation are, therefore, not entirely my own.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

The three essays presented in this dissertation study topics in international economics, macroeconomics, and the interactions between the two.

The first chapter studies the impact of globalization on the income gaps between the rich and the poor. This paper presents a new piece of empirical evidence showing that the access to global market, either through exporting or through multinational production, is associated with higher executive-to-worker pay ratio within the firm. It further shows that firm-level inequality is higher among exporting and multinational firms because those firms are on average larger than non-exporting firms.

It then builds a model with heterogeneous firms, occupational choice, and executive compensation to model analytically and assess quantitatively the impact of globalization on the income gaps between the rich and the poor. The key mechanism is that the "gains from trade" are not distributed evenly within the same firm. The compensation of an executive is positively linked to the size of the firm, while the wage paid to the workers is determined in a country-wide labor market. Any extra profit earned in the foreign markets benefits the executives more than an average worker. Consistent with the empirical patterns described above, in the model the size of the firm solely determines the level of firm-level inequality, therefore once the size is controlled for, the exporting status of a firm has no impact on its executive-to-worker pay ratio.

The model is then calibrated to create a counterfactual world where the only source of change is the access to the global markets. The model-generated top income shares closely resemble the dynamics of income shares in the U.S. data. The correlation between the model-generated income share and the data is 0.95 for the top 0.01 percent. The adjusted R-squared of regressing the data sequence against the model-generated sequence is 0.89. In terms of magnitude, the surge in top 0.01 percent income shares in the model is about 33 percent of the surge in the data.

The second chapter, joint with Rüdiger Bachmann, studies the link between the physical micro environment (frictions and heterogeneity) and the macroeconomic dynamics in general equilibrium macro models. Specifically. this chapter studies the relationship between nonconvex capital adjustment costs at the firm level and aggregate investment dynamics. We study this question quantitatively with a two-sector lumpy investment model with inventories. We find that with inventories, nonconvex capital adjustment costs dampen and propagate the reaction of investment to shocks: the initial response of fixed capital investment to productivity shocks is 50% higher with frictionless adjustment than with the calibrated capital adjustment frictions, once inventories are introduced. The reason for this result is that with two means of transferring consumption into the future, fixed capital and inventories, the tight link between aggregate saving and fixed capital investment is broken. In contrast, in the case the literature has focused on with only one type of capital good to save and invest in, fixed capital investment dynamics are more tightly linked to consumption dynamics, which, in turn, are determined by the Euler equation of a representative household, which holds regardless of whether fixed capital investment is costly or not.

The last chapter, joint with Rüdiger Bachmann and Andrei Levchenko, presents a set of novel empirical facts that the aggregate U.S. imports, exports, and real exchange rate show conditional heteroscedasticity. We estimate two ARCH family time series models and show that conditional heteroscedasticity is statistically significant for imports and exports between 1970 and 2012, and for the real exchange rate between 1973 and 2012. Furthermore, we break down aggregate imports and exports into goods and services. We find that in most specifications, the imports and exports of goods exhibit conditional heteroscedasticity, while the trade in services does not. These new empirical findings are important — they imply that the response of key international economics variables to policy interventions probably depends on the history of past shocks. Specifically, they suggest that the response of these variables might depend on whether we are at the peak or at the bottom of a business cycle. These empirical findings run against the predictions of many models in the international economics literature. For example, the international real business cycle (IRBC) model introduced in *Backus et al.* (1992) will predict a different pattern — the impact responses of imports, exports, and real exchange rates in an IRBC model only depend on the magnitude of the current shock, but not on the history of past shocks. At the end of the paper we discuss the possibility of incorporating inventory dynamics into a trade model, and the potential mechanisms that can be used to understand the new empirical findings.

CHAPTER II

Globalization and Top Income Shares

2.1 Introduction

The real income of the top 0.01 percent of the population increased by 118.5 percent between 1993 and 2011 in the United States. Over the same period, the real income of the bottom 99 percent increased by 5.8 percent only (*Piketty and Saez* (2003))[1]. What is the role of globalization in the widening income gap between the very rich and the rest of the population? Unfortunately, we do not have a good answer. Researchers working on the distributional effects of trade usually focus on wage inequality and especially on the "skill premium," which is the wage difference between the skilled and unskilled workers.[2] However, the income of the top 0.01 percent – which usually consists of executive compensation, business profits, and capital gains – cannot be easily explained using the "skill premium." For example, numerous studies have shown that education level, a widely used measure of skill, has no clear correlation with CEO compensation.[3]

Complementing the literature on the "skill premium," this paper focuses on the income gaps between the rich and the rest of the population. I first document that

[1]Data updated to 2011.

[2]Among many others, see *Goldberg and Pavcnik* (2007), *Helpman et al.* (2010), and *Burstein and Vogel* (2012).

[3]For example, see *Belliveau et al.* (1996) and *Geletkanycz et al.* (2001) for details.

within the same firm, the income gaps between the top executives and the average worker are higher among exporting and multinational firms than among non-exporting firms in the United States. To do so, I link data on executive compensation in both publicly traded and privately held firms to confidential data from the U.S. Census Bureau to create a new data set that covers firm-level executive compensation, employment, payroll, and export sales. I find that, on average, the CEO-to-worker pay ratio within the same firm is around 41 to 50 percent higher among exporting and multinational firms than among domestic firms. Similar results can be found for the income gaps between other top executives and the average workers as well. These empirical findings suggest that globalization might be responsible for the widening income gaps between the rich and the poor through within-firm inequality, a channel that has rarely been explored. Moreover, I find that the "exporter premium" in firm-level inequality is mainly driven by the size premium of exporting and multinational firms. Once the size of the firm is controlled for, the between-group differences in firm-level inequality are no longer significantly different from 0. The link between firm-level inequality and size suggests that the empirical findings can be naturally rationalized in a model where the superiority in size is associated with exporting status.

I therefore propose a new framework that bridges the heterogeneous firm trade model based on *Melitz* (2003) and *Helpman et al.* (2004) with the literature of occupational choice and executive compensation. The model world consists of two countries. Each country is populated by a fixed measure of individuals who are endowed with different levels of human capital. Individuals can choose between different occupations, as in *Lucas* (1978). They can either (1) create a new firm and become the founder and CEO of the firm or (2) work for an existing firm. If they choose to create a new firm, their human capital determines the productivity of the firm, and their income depends positively on the size of the firm they create. If they choose to be workers, their human capital determines the amount of efficiency labor they supply to the market. The wage rate of efficiency labor is determined in a competitive countrywide labor market and equalized across firms within the same country. In equilibrium, only the individuals with human capital above a certain threshold choose to create firms, while the majority of the population choose to be workers. The production and consumption sides of the economy are modeled following *Helpman et al.* (2004). Each firm produces one variety of goods in a monopolistically competitive market. Firms have two options to sell to the foreign markets: they can either export or set up subsidiaries abroad (multinational production). Individuals cannot move across borders, and they consume a constant elasticity of substitution (CES) aggregate of all the available varieties in their country.

In equilibrium, within-firm inequality is higher among the firms that sell to the foreign market. The key mechanism is that the "gains from trade" are not distributed evenly within the same firm. The compensation paid to the CEO of a firm is linked to the size of the firm, while the wage rate of a typical worker is determined in a countrywide labor market. Therefore, any extra profits earned in the foreign market directly benefit the CEO, but not the workers. In the end, as the firm starts to sell to the global markets, its within-firm inequality will be higher. On the aggregate level, this will create a gap in firm-level inequality between the exporting and domestic firms. Consistent with the empirical patterns described above, in the model, the size of the firm solely determines the level of within-firm inequality; therefore, once the size is controlled for, the exporting status of a firm has no impact on its CEO-to-worker pay ratio.

In addition to rationalizing within-firm inequality, the model also offers a parsimonious way to capture the U.S. income distribution and firm size distribution at the same time. Empirically, the U.S. income distribution is well approximated by an exponential distribution for the majority at the left tail and a fat-tailed Pareto distribution for the right tail.[4] At the same time, the U.S. firm size distribution can also be well described by a fat-tailed Pareto distribution (*Axtell* (2001)). These two distributions are captured simultaneously within the model by two assumptions: (1) human capital is distributed exponentially, and (2) firm productivity is an exponential function of the founder's human capital. The model then features a Pareto firm size distribution, while the income distribution follows a two-class structure. The individuals at the right tail of the income distribution are the CEOs, whose income is linked to the size of the firm they manage. This implies that the right tail of the income distribution will follow the firm size distribution and thus be Pareto. The workers' wage depends on their human capital, which implies an exponentially distributed income outside of the right tail. When the model is calibrated to capture the firm size distribution, it reproduces the income distribution observed in the data with reasonable precision. For example, the top 1 percent income share in 2008 in the model is 16.59 percent, while it is 17.89 percent in the data. The top 0.01 percent income share is 3.83 percent in the model, while it is 3.37 percent in the data in the same year.

The calibrated model is then used to quantitatively assess the impact of globalization on the top income gap. I calibrate the barriers to trade so that the imports-to-GDP ratio and multinational-firms-sales-to-GDP ratio in the model match their counter-parts in the U.S. data in each year between 1988 and 2008. All the other parameters are held constant, creating a counterfactual world where the only source of change is the access to the global markets. The model-generated top income shares closely resemble the trajectory of income shares in the data. The correlation between the model-generated income share and the data is 0.95 for the top 0.01 percent. The adjusted R-squared of regressing the data sequence against the model-generated sequence is 0.89. In terms of magnitude, the surge in the top 0.01 percent income

[4] See *Drăgulescu and Yakovenko* (2001a), *Drăgulescu and Yakovenko* (2001b), *Clementi and Gallegati* (2005), and *Yakovenko and Silva* (2005) for details.

shares in the model is about 33 percent of the surge in the data. The change of top income shares explained by globalization is high considering that none of the other major sources of income inequality, such as equity markets and income tax systems, is present in this model.

This paper is related to several strands of literature. First, it contributes to the literature on the distributional effects of globalization. The majority of the existing research focuses on wage inequality (see, for example, *Feenstra and Hanson* (1996), *Manasse and Turrini* (2001), *Yeaple* (2005), *Helpman et al.* (2010), and *Egger and Kreickemeier* (2012)). This paper is the first to focus on the right tail of the income distribution. This paper is also linked to the executive compensation literature (e.g. *Roberts* (1956), *Baker and Hall* (2004), *Gabaix and Landier* (2008), *Frydman and Saks* (2010)). The main mechanism of this paper (i.e., that the elasticity between CEO compensation and firm size is positive) is based on the empirical findings of this literature. This paper contributes to this literature by introducing census data to the study of executive compensation. Section 2.2 provides a detailed discussion on the advantages of using census data in this literature.

Lastly, this paper is related to the occupational choice literature going back to *Lucas* (1978). Efforts to merge occupational choice models with trade models are rare. *Monte* (2011) and *Meckl and Weigert* (2011) developed models similar to the one presented here. Their models are used to study wage distributions and the origin of firm productivity. This paper focuses on the income gaps between the rich and the poor.

The rest of this paper is organized as follows. Section 2.2 presents the empirical results. Section 2.3 discusses the model setup and Section 2.4 focuses on the analytical results. Section 2.5 provides details of the calibration process and Section 2.6 provides the quantitative results. Section 2.7 presents the conclusions.

2.2 Empirical Results

2.2.1 Evidence from U.S. Publicly-Traded Firms

2.2.1.1 Main Results

In this section, I document that the income gaps between the executives and average workers are higher for U.S. public firms engaged in the global market between 1992 and 2007.

The empirical evidence is based on a linked data set that has three components: ExecuCompustat from Standard & Poor, the Longitudinal Business Database (LBD) from the Census Bureau, and the Longitudinal Firm Trade Transactions Database (LFTTD) from the U.S. Customs and the Census Bureau. The ExecuCompustat provides data on executive compensation. It reports the total realized and estimated compensation of the CEO, CFO, and three other highly paid executives of U.S. public firms in the S&P Composite 1500 Index from 1992 onward.[5] The executive compensation consists of salary, bonus, stock options, long term incentive plans (LTIPs), restricted stock awards, and all others. "Realized" compensation (variable name: TDC2) measures the value of stock option awards at the time of execution, while "estimated" compensation (variable name: TDC1) measures the value of stock options at the time of granting using the Black-Scholes formula.[6]

The confidential Census Bureau databases provide the other key variables needed to measure within-firm inequality and exporting status. The LBD is compiled from the Census Bureau's Business Register, which covers the universe of U.S. firms at the

[5] The Securities and Exchange Commission (SEC) requires public firms to disclose the total compensation of at least five said executives starting from 1992. Any firm that was once included in the S&P 1500 Index is included in the sample, even if the firm is later dropped from the index. The S&P 1500 Index is the union of three commonly used indices: S&P 500 (LargeCap), S&P MidCap 400 Index, and S&P SmallCap 600 Index. This index covers approximately 90 percent of the total U.S. public firm capitalization.

[6] In 2006, the SEC changed the disclosure rule on executive compensation, which makes the raw data before and after 2006 not directly comparable. The ExecuCompustat data set takes this into account when constructing TDC1 and TDC2 so these two variables can be used for the entire sample.

establishment level. I aggregate it up to the firm level and extract annual employment and payroll variables, which are used to compute the average non-executive wage for each firm in a given year. The LBD is linked to the last component of the data set, the LFTTD, using the methods described in *McCallum* (2013). The LFTTD records the universe of individual international trade transactions made by U.S. firms based on the data collected by U.S. Customs from 1992 onward. It links each export transaction to the U.S. exporting firm and thus provides the base to identify exporting firms in each year. The final linkage between ExecuCompustat and the linked LBD-LFTTD is done through the Compustat-SSEL Bridge provided by the Census Bureau.

This paper is the first to use census data in the study of executive compensations. The linked data set has several advantages relative to the data used in the existing literature. The first is the coverage of employment and payroll data. U.S. public firms are not required to disclose non-executive compensations. As a result, the majority of firms do not report total payroll expenditure in SEC filings, making it almost impossible to compute wages at the firm level. For example, as reported by *Faleye et al.* (2013), around 90 percent of the firm-year observations and 87 percent of unique firms have to be dropped from ExecuCompustat due to this missing value problem in their study of the CEO-to-worker pay ratio. The under-reporting also leads to distortions of sectoral representation in the sample. Financial firms are overrepresented in the main sample of Faleya et al (2013), which are responsible for 47 percent of the firm-year observations, while they are responsible for only 15 percent of observations in ExecuCompustat. Similarly, manufacturing firms are underrepresented (16.2 percent in Faleya et al (2013) v.s. 43.0 percent in ExecuCompustat). In comparison, the LBD provides universal coverage of employment and payroll and thus minimizes the loss of observations. Table A.1 compares the number of observations and the sectoral distribution between the linked data set and the original ExecuCompustat. Overall, around 50 percent of the ExecuCompustat observations can be matched with the linked LBD-LFTTD. The sectoral distribution of firms in ExecuCompustat also compares well to the linked data set. For example, financial firms are responsible for around 14 percent and manufacturing firms 47 percent in the linked data set. The second advantage of the linked data set is the identification of exporting firms. Again, as firms are not required to report export sales separately, the missing value problem is prevalent, forcing the researcher to discard a large proportion of the data set in studies that involve exporting behavior. This issue is solved by using the LFTTD, which provides universal coverage of U.S. international transactions and thus minimizes the loss of observations.

The final linked data set contains a sample of 17,233 firm-year observations between 1992 and 2007 with 2,561 unique firms. A total of 13,169 firm-year observations are classified as exporters and the remaining 4,054 as non-exporters. The key variable of interest is the CEO-to-worker pay ratio. I construct this ratio as the total realized compensation (TDC2) divided by the average non-executive wage. I use realized compensation as the benchmark measure of CEO income instead of estimated compensation. This is because the former can be directly observed, while the latter has to be inferred from a pricing formula. I report the robustness checks using total estimated compensation (TDC1) in Appendix A.3, and the results are essentially the same. The average non-executive wage is the total payroll of a firm in a given year minus the salary and bonus of the CEO, then divided by total employment minus 1. The reason for this definition is as follows: "Total payroll," as reported in the LBD, comes from the Business Register, which is in turn based on IRS tax records. The salary and bonus of the CEO are reported as part of the total payroll for tax purposes, while the income earned from stock options is not.[7] Therefore, I only need to subtract part of the total compensation when computing the non-executive wage. The denominator is one less the total employment to account for the fact that the

[7]The "total payroll" and "employment" items in LBD are compiled from filings of IRS Form-941/943. See IRS Publications 15, 15-A, and 15-B for the details of tax deductions and exemptions.

CEO is also counted as an employee in tax filings.

As reported in Table A.2, on average, the CEO earns 89 times more than an average worker in his/her own firm across the entire sample. The CEO-to-worker pay ratio varies by exporting status: it is 92 for exporting firms and 81 for non-exporting firms. To test the difference in the CEO-to-worker pay ratio between the two groups, I estimate the following equation with the pooled panel data:

$$\log(\mathrm{CEO_{it}}/\mathrm{WAGE_{it}}) = \beta_0 + \beta_1 \mathrm{EXP}_{it} + \mathbf{b'}_2 \cdot \mathbf{s} + \mathbf{b'}_3 \cdot \mathbf{y} + \epsilon_{it}, \tag{2.1}$$

where $\mathrm{CEO_{it}}/\mathrm{WAGE_{it}}$ is the CEO-to-worker pay ratio, EXP_{it} is the exporter status indicator for firm i at year t, $\mathbf{s}$ is a vector of sector dummies at a four-digit Standard Industrial Classification (SIC) level, and $\mathbf{y}$ is a vector of year dummies. The standard errors are clustered at the year-sector level. The coefficient of interest is β_1: if the CEO-to-worker pay ratio is significantly higher for exporters, we shall expect this parameter to be positive.

The first column in Table 2.1 confirms that the "exporter premium" in firm-level inequality exists after controlling for time and sector fixed effects. The estimated β_1 is statistically significant at the 0.01 level, and the size of the coefficient suggests that the gap between the groups is large. On average, the CEO-to-worker pay ratio is 50.7 percent higher among exporters than among non-exporters, with a standard error of around 3 percent.

Why is within-firm inequality higher among exporters? The other columns of Table 2.1 try to shed some light on this. The second column controls for the size of the firm by introducing the logarithm of annual sales, as reported in Compustat into the right-hand side of equation (2.1) and the third column drops the exporting indicator but keeps the logarithm of annual sales in equation (2.1). These three columns together suggest that the "exporter premium" in within-firm inequality is driven by

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exporter	0.507***	0.0238		0.0621**		0.0700***	
	(0.0300)	(0.0246)		(0.0247)		(0.0270)	
Sales		0.436***	0.437***				
		(0.00719)	(0.00702)				
Asset				0.420***	0.425***		
				(0.00692)	(0.00676)		
Payroll						0.364***	0.370***
						(0.00864)	(0.00828)
Constant	2.017***	-0.144	-0.132	-0.0576	-0.0248	-1.574***	-1.569***
	(0.210)	(0.196)	(0.196)	(0.202)	(0.201)	(0.216)	(0.216)
Observations	17223	17223	17223	17223	17223	17223	17223
R-squared	0.270	0.439	0.439	0.428	0.428	0.385	0.385

Table 2.1: CEO-to-Worker Pay Ratio of U.S. Public Firms

Note: This table reports the results of estimating equation (2.1) for U.S. public firms based on the linked ExecuCompustat-LBD-LFTTD data. The left-hand side variable for each of the regressions is the (log of) CEO-to-worker pay ratio. “Exporter” is the exporter indicator computed from LFTTD. “Sales” is the (log of) total annual sales reported in Compustat. “Asset” is the (log of) total asset reported in Compustat. “Payroll” is the (log of) total annual payroll reported in LBD. The unit of observation is firm-year and year varies between 1992 and 2007. All specifications include year and four-digit SIC fixed effects. See Table A.1 for sector distribution of the sample. Robust standard errors are clustered at the year-sector level.

the size premium of the exporters. Comparing the first and second columns, the coefficient on the exporting indicator drops to 2 percent and is no longer significantly different from 0 once the sales of the firm is controlled for. In contrast, the comparison between the second and third columns suggests that introducing the exporter indicator on top of the size variable does not change the results significantly. The estimated coefficient on annual sales and the adjusted R-squared barely move, if at all, between these two columns. Columns 4 to 7 repeat the same exercises with other controls for the size of the firm, such as total asset, as reported in Compustat, and total U.S. payroll, as reported in the LBD.[8] When the total asset of the firm is controlled for, the “exporter premium” in inequality drops to around 6 percent. When the total payroll is controlled for, the coefficient drops to around 7 percent.

These exercises convey a consistent message. The difference in the CEO-to-worker pay ratio is mainly driven by the size difference between firms. Larger firms tend to

[8]Other controls such as total employment and combinations of different controls are reported in Table A.3 in Appendix A.3.

exhibit higher within-firm inequality, and the reason we observe higher within-firm inequality among exporters is precisely because those firms are, on average, larger – a stylized fact confirmed by the large empirical trade literature that motivated the new generation of heterogeneous firms trade models.[9] This suggests that within-firm inequality can be naturally incorporated into a Melitz trade model, where exporting behavior and size are linked.

The insignificance of exporting status conditional on size does not imply that trade is irrelevant for within-firm inequality. Without trade, many of the large firms in the sample will not be able to grow to the size that we observe in the data in the first place. In a counterfactual world where all the firms can only sell to the domestic market, many of the large firms would be smaller, and thus, their within-firm inequality would be lower. The insignificance of the exporter dummy conditional on size only implies that whatever effect trade might have on within-firm inequality, the main channel is the size of the firm. In some cases, the coefficient on exporter dummy is significantly positive after controlling for size, indicating that there are other factors that predict higher within-firm inequality among exporters. For example, exporting firms might need different managerial skills than domestic firms and thus are recruiting their CEOs in a different market. However, as the size of the coefficients suggests, no matter what these factors are, their explanatory power is small relative to firm size. Therefore, the model presented in Section 2.3 focuses solely on the size of the firm and leaves the other factors to future research.

2.2.1.2 Extensions

The benchmark result analyzes how the ratio between total CEO compensation and wage varies between firms of different exporting status. In the rest of the section, I extend the analysis in three directions: First, I decompose the total CEO

[9]For example, see *Bernard and Jensen* (1999).

compensation into three different components and study whether the same pattern can be detected in each part. I then shift the focus from the CEO to other highly paid executives, and in the last, I study whether the CEO-to-worker pay ratio varies across multinational status.

Components of Executive Compensation Executive compensation often consists of salary, bonus, stock options, and LTIP.[10] While some of the items such as stock options and bonuses are volatile and linked to the performance of the firm, other items such as salary are much less so. Is the "exporter premium" in executive-to-worker pay ratio driven by certain components? To answer this question, I decompose executive compensations into three parts: "salary," "bonus," and "stock options and others"[11] and estimate equation (2.1) for each part separately.

The results are presented in Tables A.5 and A.6. The same pattern can be observed in all three components of the CEO compensation: the CEO-to-worker pay ratio is higher among exporters, whether we measure the CEO compensation using salary, bonus, or stock options. On average, the stock-options-to-wage ratio is around 85 percent higher among exporters than among non-exporters. The bonus-to-wage ratio is 51 percent higher, while the salary-to-wage ratio is about 21 percent higher. The "exporter premium" in stock and option rewards is the highest among the three components. It could be that exporting firms usually face additional risks related to international trade such as exchange rate uncertainty and disruptions to trade routes. Part of the higher premium in stock and option rewards is probably the compensation to the higher risk. This also applies to the premiums observed in bonus, though to a lesser extent. The coefficient on salary, the riskless component of compensation, is also significantly different from 0. This implies that risk premium cannot fully explain the "exporter premium" in CEO-to-worker pay ratio. The coefficient on salary is the

[10]See *Murphy* (1999) for details.

[11]The items in the "other" category include LTIPs such as restricted stock plans and multi-year accounting-based performance plans.

smallest also because the correlation between firm size and salary is relatively weak: many large firms optimally choose to use other rewards to substitute for salary for accounting and tax purposes.[12]

Top 5 Highest Paid Executives Does the gap in the CEO-to-worker pay ratio between exporters and non-exporters extend to other top executives? The data allow me to expand the focus of executive compensation to include the top five highest paid executives in each firm. I define the "top five highest paid executives" by their total realized compensation. Robustness checks using total estimated compensation are provided in Appendix A.3. The CEO of each firm is always included in the sample, and the CFO is included in most cases. I measure within-firm inequality using the ratio between average compensation of the top 5 executives and the non-executive wage. To compute non-executive wage in this case, I subtract the salary and bonus of all five executives from the total payroll, and then divide it by total employment minus 5.

The estimation of equation (2.1) is reported in Table A.4. The main results carry over when we include the other highly paid executives. On average, the top-5-to-worker pay ratio is 47 percent higher among exporters than among non-exporters. Once again, the "exporter premium" of inequality is driven by the size premium: once the size of the firm is controlled for, either by sales, asset, employment, or payroll, the estimated coefficient on exporter indicator drops by a large margin to around 0 to 7 percent.

Multinational Firms The last extension introduces multinational firm indicators. The multinational firm indicators are constructed from the geographic segment data

[12]For example, the provisions to the 1993 legislation "Omnibus Budget Reconciliation Act of 1993" put a $1 million cap on the deductibility of "non-performance based" executive compensations (the so-called Section-162 $1 million rule). This rule primarily reduces the incentives for large firms to pay high salaries but has a limited effect on bonus, stock options and, total compensations in general. See *Rose and Wolfram* (2002) for details.

in Compustat. I classify a firm-year observation as multinational if a U.S. firm reports the existence of a non-domestic geographic segment. I then screen all the names of the segments and divisions to correct for apparent errors in the geographic location indicator. The multinational indicators from segment data are then linked with the ExecuCompustat-LBD. The resulting data set contains 12,943 firm-year observations and 1,606 unique firms. Out of these firm-year observations, 5,885 records are classified as non-MNE and the rest 7,058 as MNE. On average, the CEO-to-worker pay ratio is 87.4 among the non-MNE group and 100.0 among the MNE group.

I re-estimate equation (2.1) with the MNE data, and the results are reported in Table A.7. On average, the CEO-to-worker pay ratio is 28.4 percent higher among the MNE group than among the non-MNE group. After controlling for annual sales and assets, the between-group difference is no longer significantly different from 0. After controlling for the employment and total payroll of their U.S. operations, the MNE group sees around 14.0 percent higher in the CEO-to-worker pay ratio. The difference between controlling for sales/total asset and employment/payroll is that the former group is based on the aggregate of global operations, while the latter group is based on U.S. operations only. Nevertheless, under all four controls, the size of the "MNE premium" in within-firm inequality drops significantly, indicating that the size premium of multinational firms can explain the majority of the difference in within-firm inequality across the groups.

2.2.2 Evidence from U.S. Privately Held Firms

The results in the previous section cover publicly traded firms. In this section I document that the same empirical pattern can be observed among privately owned firms in the U.S. The majority of U.S. firms are private, and they are responsible for more than 60 percent of firm sales in 2007.[13] The executive compensation is believed

[13]Total sales of U.S. firms come from the Census Bureau's *Statistics of U.S. Businesses, 2007*. Total sales of public firms come from Compustat.

to be less affected by corporate governance problems in private firms because they tend to be more closely held.[14] These considerations make the private firms particularly interesting subjects to examine in this paper.

Unlike public firms, most private firms are not subject to SEC's executive compensation disclosure rules. As a result, the majority of data sets on private firm compensations have to collect data through surveys. Survey-based data sets usually do not disclose firm identifiers, such as names and addresses, which makes the link to the census data impossible.

The data set used in this paper comes from the Standard & Poor's Capital IQ (CIQ). Unlike the survey-based data set, CIQ collects data through regulatory filings,[15] news aggregators, and company websites. The "People Intelligence" part of the data set covers over 4.5 million professionals across the globe, many of them working as executives in privately held companies. The unique advantage of CIQ data is that they provide the names, addresses, and telephone numbers for all the firms covered, making the linkage to the census data possible.

I start with executives working in private U.S. firms between 2003 and 2007 from the CIQ data. This yields a data set that contains around 33,000 individuals working in 3,849 privately held firms and 11,706 firm-year level observations. I then link this data set to the Standard Statistical Establishment List (SSEL) in the Census Bureau. Unlike the ExecuCompustat, where the bridge files exist and firms can be matched using standardized identifiers, the CIQ data have not been linked to the census data sets before. Therefore, I carry out a fuzzy match based on name, street address, and zip code. I require that the weighted similarity has to be at least 95 percent for two entries to be considered a match and then hand-screen all the matched records to eliminate obvious errors. The matched CIQ records are then linked with

[14] See *Jensen* (1997) and *Hartzell and Starks* (2003) for details.

[15] Contrary to common belief, some private firms are subject to executive disclosure rules similar to public firms by the SEC. They are usually large firms with more than 500 shareholders and more than \$10 million in total assets. See *Gao et al.* (2012) for more details.

LBD-LFTTD constructed by *McCallum* (2013), creating, to my knowledge, the first data set that contains private firm executive compensation and reliable measures of exporting status, employment, and payroll at the firm level.

Table A.8 summarizes the results of the fuzzy merge and compares the distribution of firms across sectors in the linked data set and the original CIQ data. The linked data set contains 6,002 firm-year observations and 2202 unique firms. A total of 3,366 firm-year observations and 1,207 unique firms are exporting firms, while the remaining 2,636 observations with 9,95 unique firms are non-exporters. Overall, 51 percent of the CIQ records are matched with the census data. The sectoral distribution of the CIQ is preserved in the linked data set. For example, manufacturing firms constitute 33.8 percent in the linked data and 34.4 percent in the original data; financial firms are responsible for 22.0 percent in the linked data and 18.9 percent in CIQ.

Instead of the CEO-to-worker pay ratio, I construct the ratio between the highest-paid executive and the non-executive wage as the benchmark measure of intra-firm inequality. The CIQ data does not report standardized job titles, and therefore, constructing the CEO title from the raw data would introduce unnecessary noise. Nevertheless, most of the highest-paid executives are indeed CEOs: in ExecuCompustat, more than 98 percent of the highest-paid executives are the CEOs. There is no strong reason to believe that this ratio will be significantly different in the CIQ sample.

The summary statistics of the top-1-to-worker pay ratio are reported in Table A.9. Overall, within-firm inequality is lower among private firms than among public firms. The top-1-to-worker pay ratio is 37.6 in the private firm sample compared with 89 in the public firm sample. Again, the top-1-to-worker pay ratio varies with exporting status. The ratio is 41.3 among exporters and only 32.8 among non-exporters.

The results of re-estimating equation (2.1) using the linked CIQ data are reported in Table 2.2. Overall, the results are similar to those based on the public firm sample.

On average, the top-1-to-worker pay ratio is 41 percent higher among exporters than among non-exporters. This difference is statistically different from 0, with a standard error of 5.4 percent. Again, the gap between exporters and non-exporters is mainly driven by the size difference. Once I control for total sales, total asset, or total payroll, the estimated coefficient on exporting status is no longer significantly positive. In the case of controlling for total payroll, the estimated "exporter premium" is even slightly negative.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exporter	0.411***	-0.0630		-0.0441		-0.0932**	
	(0.0546)	(0.0415)		(0.0419)		(0.0426)	
Sales		0.388***	0.384***				
		(0.0108)	(0.0105)				
Asset				0.391***	0.389***		
				(0.00910)	(0.00916)		
Payroll						0.353***	0.345***
						(0.0112)	(0.0105)
Constant	2.690***	0.844***	0.831***	0.792***	0.782***	-3.248***	-3.169***
	(0.199)	(0.233)	(0.230)	(0.221)	(0.219)	(0.307)	(0.300)
Observations	6002	6002	6002	6002	6002	6002	6002
R-squared	0.402	0.619	0.618	0.627	0.627	0.566	0.565

Table 2.2: Highest-Paid-Executive-to-Worker Pay Ratio of U.S. Private Firms

Note: This table reports the results of estimating equation (2.1) for U.S. private firms based on the linked CIQ-LBD-LFTTD data. The left-hand side variable for each of the regressions is the (log of) highest-paid-executive-to-worker pay ratio. "Exporter" is the exporter indicator computed from LFTTD. "Sales" is the (log of) total annual sales reported in CIQ. "Asset" is the (log of) total asset reported in CIQ. "Payroll" is the (log of) total annual payroll reported in LBD. The unit of observation is firm-year and year varies between 2003 and 2007. All specifications include year and four-digit SIC fixed effects. See Table A.8 for sector distribution of the sample. Robust standard errors are clustered at the year-sector level.

I repeat the first two extensions in the previous section.[16] I first extend the analysis to include all top 5 highly paid executives. The results are presented in Table A.11 and are again similar to those obtained in the public firm sample. On average, the top-5-to-worker pay ratio is 39 percent higher among exporters than among non-exporters, and the gap is mainly driven by size differences in the firms: once the size of the firm is controlled for, the estimated coefficient drops to between -1.4 and -7.7

[16]The private data set does not contain enough information to identify multinational firms; therefore, the last extension can not be repeated here.

percent.

I also decompose the total realized compensation into three parts (i.e., salary, bonus and all others) and re-estimate equation (2.1).[17] The results are presented in Tables A.12 and A.13. The "exporter premium" in within-firm inequality exists in all three components, with the same ranking of magnitude as in the public firm sample. The "all-others"-to-wage ratio is 54 percent higher among exporters than non-exporters, followed by the bonus-to-wage ratio (31 percent) and the salary-to-wage ratio (16 percent). The size of the "exporter premium" drops significantly for all three components once I control for the size of the firm.

The estimation results from the public and private firms in the U.S. convey the same message. Firm-level inequality, measured by various executive-to-worker pay ratios, is higher among exporting firms than among non-exporting firms. The gaps in firm-level inequality between exporters and non-exporters are driven by the size difference between the groups. This finding suggests that the observed "exporter premium" in within-firm inequality should be incorporated into a heterogeneous firm trade model, where the superiority in size is associated with exporting behavior.

2.3 The Model

2.3.1 Environment

The model world consists of two countries indexed by i. Each country i is populated by individuals with measure n_i. People in each country are endowed with human capital x. The distribution of human capital in each country follows an exponential distribution with shape parameter λ. The cumulative distribution function (CDF) of

[17] "All others" are mainly stock and option rewards. Although these companies are not publicly traded, stock and option rewards are still popular among executives. These stocks and options are usually exercised at the time of buyout or initial public offering (IPO).

human capital is as follows:

$$F(x) = 1 - e^{-\lambda x}.$$

People can choose between two careers. They can either work for an existing firm or create a new firm. If they choose to be a worker, their human capital directly translates into the amount of efficiency labor that can be supplied to the market. In this case, their income will be $w_i x$, where w_i is the prevailing wage rate per efficiency unit of labor in country i. Individuals cannot move between the countries and the wage rate w_i is determined in a country-wide competitive labor market.

If the individual choose to create a new firm, he/she become the founder and CEO of the firm. The productivity of the firm, denoted by $A_i(x)$, depends on the human capital of the founder. Specifically, $A_i(x)$ takes the following form:

$$A_i(x) = b_i e^x, \tag{2.2}$$

where b_i is the total factor productivity (TFP) in country i. Appendix A.1.3.1 proves that A_i follows a Type-I Pareto distribution with location parameter b_i and shape parameter λ. The payoff to the founder and CEO of the firm is an increasing function of the profit of the firm, denoted as $k(\pi)$, where π is the profit of the firm. All the analytical results in Section 2.4 only require $k(\pi)$ to be monotonically increasing in π and regularly varying. For concreteness, the reader can safely assume the simplest monotonic and regularly varying function: $k(\pi) = \pi$. A more general functional form of $k(\pi)$ is assumed in the calibration of the model and details are provided in Section 2.5. For simplicity I assume that the residual profit after the CEO compensation is distributed back to the entire population in country i evenly (i.e. all the people in the country own the firms through a market mutual fund).

The production side of the economy is modeled after *Helpman et al.* (2004). A firm with productivity $A_i(x)$ produces a single variety of good, indexed by x, with the following production function:

$$q_i(x) = A_i(x) \cdot (L_i(x) - f_{ii}),$$

where $L_i(x)$ is the amount of efficiency labor used and f_{ii} is the fixed cost of production, paid in the units of labor of country i. Each firm operates in a monopolistically-competitive market and earns positive profit in equilibrium.

Firms in country i can serve the foreign market j in two ways: they can either export to country j its good produced in country i, or set up production facilities in country j and supply the market with local production (foreign direct investment, FDI). If a firm in country i wants to export to country j, it first needs to pay a fixed cost f_{ji} denominated in labor to set up the distribution network. Then trade incurs an iceberg cost of $\tau_{ji} > 1$: in order to supply country j with one unit of good from country i, the firm needs to ship τ_{ji} units out. In order to serve country j from country i through FDI, the firm needs to pay the fixed overhead costs g_{ji} in units of labor in country i. The labor costs are interpreted as the overhead costs of starting operation, as well as the costs introduced by policy barriers.

Individuals in country i consume a CES aggregate of all the varieties available in country i. Their utility function is as follows:

$$U_i = \left(\int\limits_{m \in \Theta_i} q_i(m)^{\frac{\epsilon-1}{\epsilon}} dm \right)^{\frac{\epsilon}{\epsilon-1}},$$

where ϵ is the elasticity of substitution, and Θ_i is the set of goods that are available in country i. It potentially has three subsets: (1) the goods produced by all the firms founded in country i, (2) the goods produced in country j and exported to country i, and (3) the goods produced in country i by the subsidiaries owned by the firms created in country j.

2.3.2 Solution and Equilibrium Conditions

The solution of the occupational choice problem is a single cutoff rule. There exists a human capital level x_i^* in country i such that all the individuals with human capital smaller than x_i^* choose to be workers and all the other individuals choose to create firms. x_i^* is the solution to the following equation:

$$k(\pi(x_i^*)) = w_i x_i^*. \tag{2.3}$$

Equation (2.3) requires that in equilibrium the founder of the marginal firm must be indifferent between creating a new firm or working for an existing firm. The sufficient and necessary condition for the existence and uniqueness of the solution is that $k(\pi_i(0)) < 0$, which means that the individual with the least amount of human capital must find creating a new firm unprofitable.

Figure 2.1 presents the solution in a simple setting where $k(\pi) = \pi$. The black solid line is the income of a worker as a function of his/her human capital. The blue dashed line is the income of a CEO as a function of his/her human capital. Under the assumption that $k(\pi(0)) < 0$, the two curves cross once and only once at the cutoff human capital level x_i^*.

The solution to the firm's problem is similar to *Helpman et al.* (2004). Denote the total spending in country i as H_i and the ideal price index as P_i. The maximum profit a firm in country i can earn in its domestic market is as follows:

$$\pi_{ii}(x) = \frac{H_i}{\epsilon}\left[\frac{\epsilon-1}{\epsilon}\frac{P_i}{w_i}\right]^{\epsilon-1} A_i(x)^{\epsilon-1} - f_{ii}w_i.$$

The *additional* profit a firm in country i can earn from exporting to country j is as

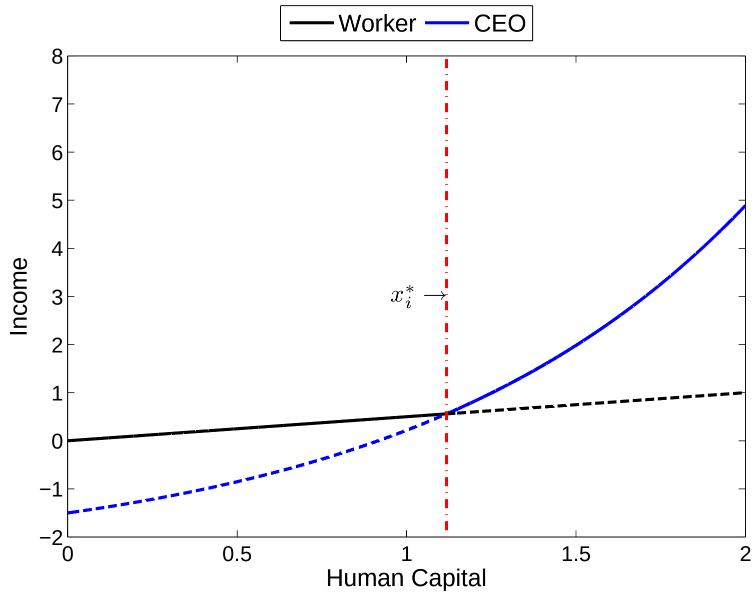

Figure 2.1: Solution of the Occupational Choice Problem

The graph plots the solution of the occupational choice problem. The black solid line is the income of a worker, and the blue dashed line is the income of a CEO. The vertical line indicates the cutoff human capital that is indifferent between being a worker or a CEO. This graph assumes that $k(\pi) = \pi$.

follows:

$$\pi^e_{ji}(x) = \frac{H_j}{\epsilon}\left[\frac{\epsilon-1}{\epsilon}\frac{P_j}{\tau_{ji}w_i}\right]^{\epsilon-1} A_i(x)^{\epsilon-1} - f_{ji}w_i, \tag{2.4}$$

and the *additional* profit a firm in country i can earn from FDI in country j is:

$$\pi^f_{ji}(x) = \frac{H_j}{\epsilon}\left[\frac{\epsilon-1}{\epsilon}\frac{P_j}{w_j}\right]^{\epsilon-1} A_i(x)^{\epsilon-1} - g_{ji}w_i. \tag{2.5}$$

The details of the solution to the firm's problem can be found in Appendix A.1.1.

Similar to the firms in *Helpman et al.* (2004), all the firms founded in country i will serve the domestic market first. Moreover, the least productive firms will only serve the domestic market. The more productive firms will serve the domestic market and the foreign market through export. The most productive firms will serve the domestic market and the foreign market through FDI. Denote the human capital of the founder of the least productive exporting firm in country i as x^e_{ji} and the human capital of the least productive MNE in country i as x^f_{ji}. These two cutoffs must be the solution to the following two equations respectively:

$$\pi^e_{ji}(x^e_{ji}) = 0, \tag{2.6}$$

$$\pi^e_{ji}(x^f_{ji}) = \pi^f_{ji}(x^f_{ji}). \tag{2.7}$$

The first condition means that the marginal exporter will find the additional profit from exporting to be 0. The second condition says that the marginal MNE shall find the profit of serving the foreign market by FDI and by exporting to be equal.

The equilibrium of the world economy is a vector of wages, $\{w_i\}$, a vector of the occupational choice cutoffs $\{x^*_i\}$, a vector of exporting cutoffs $\{x^e_{ji}\}$, a vector of FDI cutoffs $\{x^f_{ji}\}$, a vector of ideal price levels $\{P_i\}$, and a vector of total expenditures $\{H_i\}$ such that for $i = 1, 2$ and $j = 1, 2$:

1. Every individual in country i maximizes their income by solving the occupational choice problem (equation (2.3) holds).

2. Every firm optimally chooses to be a non-exporter, exporter, or multinational firm (equations (4.2) and (2.7) hold).

3. Total income equals to total expenditure in each country:

$$H_i = n_i w_i \int_0^{x^*_i} x f_i(x) dx + n_i \int_{x^*_i}^{\infty} \pi_i(x) f_i(x) dx. \tag{2.8}$$

4. Aggregate price level and the individual prices satisfy the rational expectation

condition:

$$P_i = \left(\int\limits_{m \in \Theta_i} p(m)^{1-\epsilon} dm \right)^{\frac{1}{1-\epsilon}}. \tag{2.9}$$

5. Labor market clears in each country.

Equation (2.8) is the income-expenditure identity in country i. In equilibrium, the total expenditure in country i must equal the total income in country i, which is the sum of all the wage and profit income[18]. Equation (2.9) is the definition of the ideal price index, which is the cost of one unit of utility in country i. Appendix A.1.2 provides the details on these two equilibrium conditions.

The labor market clearing condition in country i requires that total supply of efficiency labor equals to total demand. Total supply is straightforward: it equals the integral of x from 0 to x_i^* over the density function $f(x)$ (see equation (A.6) in Appendix A.1.1). Total labor demand is more complicated. It has four parts:

1. The labor used in the production of all the goods supplied to the home market i and exported to the foreign market j by the firms founded in country i:

$$L_i^{(1)} = n_i \sum_{j=1}^{2} \int\limits_{x_{ji}}^{x_{ji}^f} \frac{H_j}{P_j^{1-\epsilon}} \left(\frac{\epsilon}{\epsilon - 1} \frac{\tau_{ji} w_i}{A_i(x)} \right)^{-\epsilon} \frac{\tau_{ji}}{A_i(x)} f(x) dx.$$

2. The labor used in fixed costs of operation and export incurred for the production

[18] The CEO compensation function does not enter the total income function, because the difference between profit and CEO compensation at a given firm will be distributed back to the individuals in country i, which implies that we only need to consider total profit when accounting for total income in a given country.

in part 1:

$$L_i^{(2)} = n_i \sum_{j=1}^{n} f_{ji} \int_{x_{ji}}^{x_{ji}^f} f(x)dx.$$

3. The labor used in fixed costs for the goods supplied to country j through FDI by the firms created in country i:

$$L_i^{(3)} = n_i \sum_{j=1}^{2} g_{ji} \int_{x_{ji}^f}^{\infty} f(x)dx.$$

4. The labor used in the production of the goods supplied to country i by the foreign subsidiaries in country i from the firms founded in country j:

$$L_i^{(4)} = \sum_{j=1}^{2} n_j \int_{x_{ij}^f}^{\infty} \frac{H_i}{P_i^{1-\epsilon}} \left(\frac{\epsilon}{\epsilon - 1} \frac{w_i}{A_i(x)} \right)^{-\epsilon} \frac{1}{A_i(x)} f(x)dx.$$

2.4 Analytical Results

2.4.1 Firm Size Distribution and Income Distribution

The firm productivity distribution in country i follows a Type-I Pareto distribution with shape parameter λ and location parameter b_i. Firm sales is a linear function of $A^{\epsilon-1}$ and therefore follows Type-I Pareto distributions with shape parameter $\lambda/(\epsilon-1)$. As noted in *di Giovanni et al.* (2011), international trade systematically changes the size distribution of firms. In my framework, this is reflected in the location parameters of the sales distributions. The sales distribution for non-exporting firms has the smallest location parameter, followed by the exporting firms. The sales distribution for multinational firms has the largest location parameter.

Firm employment and profit are affine functions of $A^{\epsilon-1}$ due to the fixed costs of operation, export, and FDI. They follow Type-II Pareto distributions with shape parameter $\lambda/(\epsilon - 1)$ and two location parameters. Similar to the sales distribution, location parameters vary by the market size accessible to the firm. Appendix A.1.3 provides details on the distribution of productivity, sales, employment, and profit for different groups of firms.

Individual income is ranked by occupation: as a group, the workers earn the lowest income, followed by the domestic firm CEOs, and then followed by the CEOs at the exporting firms. The CEOs at the multinational firms occupy the pinnacle of the income distribution.

The entire income distribution follows a two-class structure. All the workers earn the same wage rate per efficiency labor unit; therefore, their income distribution is exponential, which is the same as their human capital distribution, with a different shape parameter λ/w_i.

The income of the CEOs at the top of the income distribution depends on the CEO compensation function. As stated in Section 2.3, the compensation function $k(\pi)$ is monotonically increasing in π and regularly varying. This means that for any $z > 0$:

$$\lim_{\pi \to \infty} \frac{k(z\pi)}{k(\pi)} = z^{\beta},$$

where β is called the tail index of $k(\pi)$. Appendix A.1.4 proves that under these two assumptions, the income distributions of the CEOs adopt the following CDF:

$$U(y) = 1 - y^{-\frac{\lambda}{\beta(\epsilon-1)}} R(y), y > 0,$$

where y is the income, $\frac{\lambda}{\beta(\epsilon-1)}$ is the shape parameter of the distribution, and $R(y)$ is a slowly varying function. Distributions with this form of CDF are Pareto-Type distributions and exhibit fat-tail behavior at the right end similarly as the more commonly known Type-I Pareto distribution.

Conditional on a functional form of $k(\pi)$, the location parameters of the CEO income distributions differ by the mode of access to the foreign market. The CEOs at domestic firms have the lowest location parameters. The CEOs at exporting firms follow with higher location parameters, and the CEOs at multinational firms have the highest location parameters. Appendix A.1.4 provides the detailed discussions and proofs for the workers' income distribution and various types of CEO income distributions.

2.4.2 Partial Equilibrium

The main mechanism of the model is most clearly demonstrated in partial equilibrium. Suppose that the wage rate is fixed at the labor market clearing level and that all the prices and the total expenditure are also fixed at their equilibrium values. What will happen if the two countries open up to trade?

Figure 2.2 presents the partial equilibrium results in a simplified model where FDI is shut down and $k(\pi) = \pi$. The black solid line and the blue dashed line are the same as in Figure 2.1: they are the income of workers and CEOs in autarky in the home country. When the world opens up to trade, only the most productive firms export. In the graph, the right end of the CEO income function tilts up into the red dotted line, which is the income of CEOs at those exporting firms. The gap between the red dotted line and the blue dashed line is the extra profit (and extra compensation to the CEO) earned in the foreign country. In this simple case, all the benefits of globalization, as represented by the shaded area between the two CEO income functions, are claimed by the CEOs at the exporting firms, and none of the benefits are claimed by the workers working in those firms. On the aggregate level, top income shares will be higher because the CEOs at the exporting firms are already the richest people in autarky.

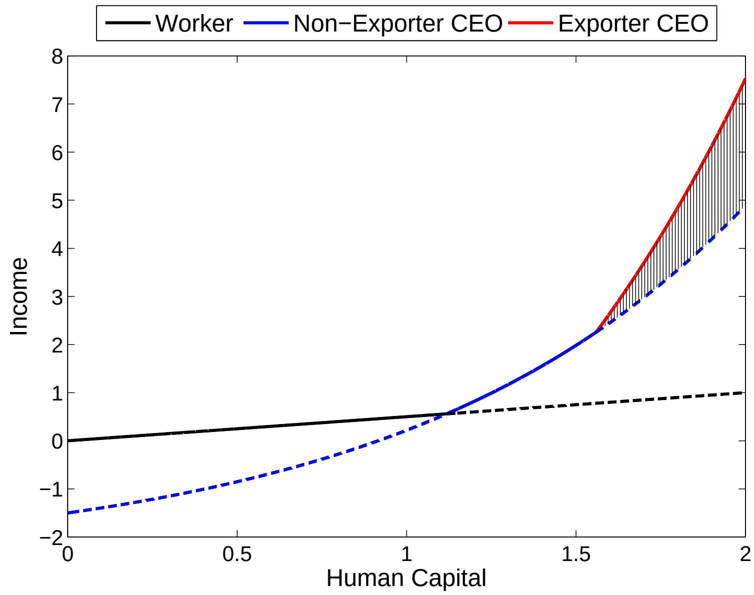

Figure 2.2: Trade and Top Income Shares in Partial Equilibrium

This graph plots the income of different individuals against their human capital for different occupations under autarky and under trade. The black solid line is the income of a worker. The blue dashed line is the income of CEOs at non-exporting firms. The red dotted line is the income of CEOs at exporting firms. The shaded area is the extra profit earned from exporting. This partial equilibrium assumes that $k(\pi) = \pi$ and that wage, total expenditure, and prices are all fixed. It also abstracts away from FDI.

2.4.3 General Equilibrium

In general equilibrium wage, total expenditure and the ideal price level respond to the changes in τ_{ij} and g_{ij}, making the results not as clear-cut as in the partial equilibrium. Nevertheless, the main mechanism of the model works the same way: the access to foreign markets benefits CEOs more than average workers, widening income inequality both at the firm and aggregate levels.

Before turning to how income inequality responds to the changes in trade barriers, I first present a simple result characterizing the cross-sectional intra-firm inequality of the model:

Proposition II.1. *If the sets of exporting firms and multinational firms in country i are non-empty, then the average CEO-to-worker pay ratio among domestic firms is strictly smaller than the average CEO-to-worker pay ratio among exporting firms,* *which in turn is strictly smaller than the average CEO-to-worker pay ratio among multinational firms.*

Proof. The least productive CEOs manage the domestic firms, which implies that, on average, they receive the least compensations among all the CEOs. The more productive CEOs manage the exporting firms, and the most productive CEOs manage the multinational firms. Since wage is equalized across the firms, the ranking of the CEO-to-worker pay ratio is the same as the ranking of the CEO income. □

Proposition II.1 implies that the empirical findings in Section 2.2 can be replicated in general equilibrium. If an econometrician observes the model world and estimates equation (2.1) without any size control, he/she will find that the CEO-to-worker pay ratio is significantly higher among firms that sell to the foreign market than those who do not. In addition, in general equilibrium, the CEO-to-worker pay ratio is proportional to any size measure of the firm. Therefore, if the econometrician can also observe any size measure of the firm and controls for it when estimating equation (2.1), the observed between-group difference will disappear, just the same as we observed in the U.S. data.

Now I turn to the results on how inequalities respond to the changes in trade. I show that the firm profit works the same way as in *Helpman et al.* (2004): as the barriers to trade get lower, the domestic firms see their profits get lower, while the high productivity exporters and multinational firms see their profits get higher. I extend these results and show that the profit-to-wage ratios in different markets follow a similar pattern. These results are summarized by the following three propositions:

Proposition II.2. *The domestic-profit-to-wage ratio, defined as*

$$\frac{\pi_{ii}(x)}{w_i} = \frac{H_i}{w_i \epsilon} \left(\frac{P_i}{w_i} \frac{\epsilon - 1}{\epsilon} \right)^{\epsilon - 1} A_i(x)^{\epsilon - 1} - f_{ii},$$

will be lower when x_i^* is higher.

Proof. See Appendix A.1.5. □

As τ_{ji} decreases, x_i^* will increase, because the marginal firm under the old τ_{ji} will no longer be profitable. This implies that the domestic profit-to-wage ratio shall be lower for all domestic firms as the country is more exposed to the global market. For the profits earned in the exporting market:

Proposition II.3. *The exporting-profit-to-wage ratio, defined as*

$$\frac{\pi_{ji}^e(x)}{w_i} = \frac{H_j}{w_i \epsilon} \left(\frac{P_j}{\tau_{ji} w_i} \frac{\epsilon - 1}{\epsilon} \right)^{\epsilon - 1} A_i(x)^{\epsilon - 1} - f_{ji},$$

will be lower when x_{ji}^e is higher.

Proof. See Appendix A.1.5. □

As τ_{ji} decreases, x_{ji}^e will be lower because the domestic firm with slightly lower productivity below the marginal exporter will now find it profitable to export. This implies that the exporting-profit-to-wage ratio will be higher for all the firms that export. A similar result for the FDI-profit-to-wage ratio obtains:

Proposition II.4. *The FDI-profit-to-wage ratio, defined as:*

$$\frac{\pi_{ji}^f(x)}{w_i} = \frac{H_j}{w_i \epsilon} \left(\frac{P_j}{w_j} \frac{\epsilon - 1}{\epsilon} \right)^{\epsilon - 1} A_i(x)^{\epsilon - 1} - g_{ji},$$

will be lower when x_{ji}^f is higher.

Proof. See Appendix A.1.5. □

Similarly, as g_{ji} gets lower, x_{ji}^f will be lower, and thus, FDI-profit-to-wage ratio will be higher.

Propositions II.3 and II.4 are the key mechanisms behind most of the quantitative results. Together, they imply that as the firm gains better access to the foreign market, the "gains from trade" are not distributed evenly within the same firm. Those whose income is linked to the profit of the firm (the CEOs) will see their income increase much faster than those whose income is not. On the aggregate level, this implies that the top income shares shall be positively linked to the volume of trade and FDI sales.

I now move on to the quantitative assessment of this mechanism. To do this, I first calibrate the model parameters in the next section.

2.5 Calibration

I interpret the two countries in the model world as the U.S. and the rest-of-the-world (ROW). I treat 109 economies combined as the ROW. These countries, together with the U.S., are responsible for around 74 percent of the world population and 82 percent of the world GDP in 2008. The list of countries included in the ROW can be found in Table A.15.[19]

The country TFP, b_i, is calculated as the Solow residual in 1988 using the methods outlined in *Caselli* (2005) and is normalized so that TFP in the U.S. is 1. The measure of population, n_i, is computed as a by-product in the estimation of the Solow residual. I first compute the "quality-adjusted workforce," as in *Caselli* (2005), using the Penn World Table 7.0 and the educational attainment data from *Barro and Lee* (2010). I then augment this measure of total workforce with the estimated capital stock and arrive at the final measure of the size of "population."[20] This measure of population takes into consideration that worker productivity varies greatly across countries because the human capital embodied in and physical capital associated with each worker varies. Using this measure, the relative size of the economies is replicated

[19]A country is included in the sample if and only if its data from 1988 to 2008 are not missing in Penn World Table 7.0 and *Barro and Lee* (2010).

[20]This "population measure" is essentially the ratio between real GDP and the estimated Solow residual in each year.

within a reasonable error margin: from 1988 and 2008, the ROW is on average 3.16 times larger than the U.S. in the data and 3.85 times larger in the model. For the details of calibrating TFP and population measures, see Appendix A.2.

The elasticity of substitution is set to be 4 so that the average markup for the firms is 33 percent. This level of mark-up is in the middle of plausible estimates.[21] The shape parameter of the human capital distribution, λ, is set to 3.81. This implies that the Pareto shape parameter of the firm employment distribution is $\lambda/(\epsilon - 1) = 1.27$, an estimation based on the LBD in 2007 for all firms with more than 10 employees.[22]

The fixed costs of operation and export are calibrated using the Doing Business database from the World Bank. Denote the days of starting a new business in each country i among the 110 countries as ϕ_i. I use the days of starting a business in the U.S. as the measure of the fixed cost of operation in the home country ($f_{11} = \phi_{\text{USA}}$). The fixed cost of operation in the ROW(f_{22}) is the GDP-weighted-average of starting a business in each of the 109 countries:

$$f_{22} = \frac{\sum_{i=1}^{109} \text{GDP}_i \phi_i}{\sum_{i=1}^{109} \text{GDP}_i}.$$

The fixed costs of export are calibrated as follows: I first compute, among the 110 countries, the fixed costs of export from country i to country j as the sum of the days of starting a business in these two countries. Denote the sum as ϕ_{ij}:

$$\phi_{ij} = \phi_i + \phi_j.$$

The fixed cost of exporting from the U.S. to the ROW (f_{21}) is computed as the

[21]For example, *Domowitz et al.* (1988) estimated the average markup for U.S. manufacturing firms to be 0.37. *Rotemberg and Woodford* (1991) used steady-state markups between 0.2 and 0.6, while *Feenstra and Weinstein* (2010) estimated the average markup to be 0.3 in 2005 in the U.S. The elasticity of substitution used here is slightly lower than the estimates based on gravity equations, which are usually between 5 and 10, as reported by *Anderson and van Wincoop* (2004).

[22]The shape parameter is estimated by the method of moments.

weighted average of the fixed cost of export from the U.S. to each of the 109 countries:

$$f_{21} = \frac{\sum_{i=1}^{109} E_{\mathrm{i,US}}\phi_{i,US}}{\sum_{i=1}^{109} E_{\mathrm{i,US}}},$$

where the weight, $E_{\mathrm{i,US}}$, is the export from the U.S. to each of the 109 countries. Similarly, the fixed cost of export from the ROW to the U.S.(f_{12}) is the weighted average of the fixed cost of export from each of the 109 countries to the U.S.:

$$f_{12} = \frac{\sum_{i=1}^{109} E_{\mathrm{US,i}}\phi_{US,i}}{\sum_{i=1}^{109} E_{\mathrm{US,i}}},$$

where the weight, $E_{\mathrm{US,i}}$, is the export from each of the 109 counties to the U.S. At the end, the entire f_{ij} matrix is scaled so that the exporting and multinational firms are responsible for 42 percent of the total employment in the U.S. in year 2000, as reported by *Bernard et al.* (2009).[23]

I use the following functional form of $k(\pi)$ as CEO compensation:

$$k(\pi) = \begin{cases} \pi & \text{if} \quad \pi \leq \alpha \\ \alpha^{1-\beta}\pi^{\beta} & \text{if} \quad \pi > \alpha \end{cases}, \tag{2.10}$$

This function is monotonically increasing in π and regularly varying; therefore, all the analytical results in Section 2.4 carry over. Intuitively, the function means that firms with profit less than or equal to α are "sole proprietorship" firms: the founder and CEO owns the firm and claims all the profit. Firms larger than α in profit are "corporations" not solely owned by the founder. The share of profit claimed by the CEO depends positively on the size of the firm, but the size elasticity of income (which is also the tail index of $k(\pi)$), β, is smaller than 1. Figure 2.3 plots this function based on the calibrated values of α and β.

[23]Year 2000 in the model means when the trade barrier matrices and TFP are calibrated to match the data moments in year 2000. This also applies to the calibration of α.

The function in equation (2.10) is based on the empirical findings in the literature that CEO compensation is proportional to the power function of the firm size, $k \sim \pi^{\beta}$, which is otherwise known as "Roberts law"(*Roberts* (1956), *Gabaix and Landier* (2008)). This functional form is a special case of the duo-scaling equation in *Gabaix and Landier* (2008), where α is the size of the reference firm. In this case, the reference firm is the smallest corporation in each country. I assume and check to make sure that the smallest firm in the model is always smaller than the calibrated α. This assumption ensures that both types of firms exist in equilibrium. The power function is also chosen for analytical simplicity as it is always regularly varying. Regardless of firm size, the income of the CEOs follows Pareto distributions, though the shape parameters of the CEOs at proprietorship firms and corporations differ (see Section 2.4 for details).

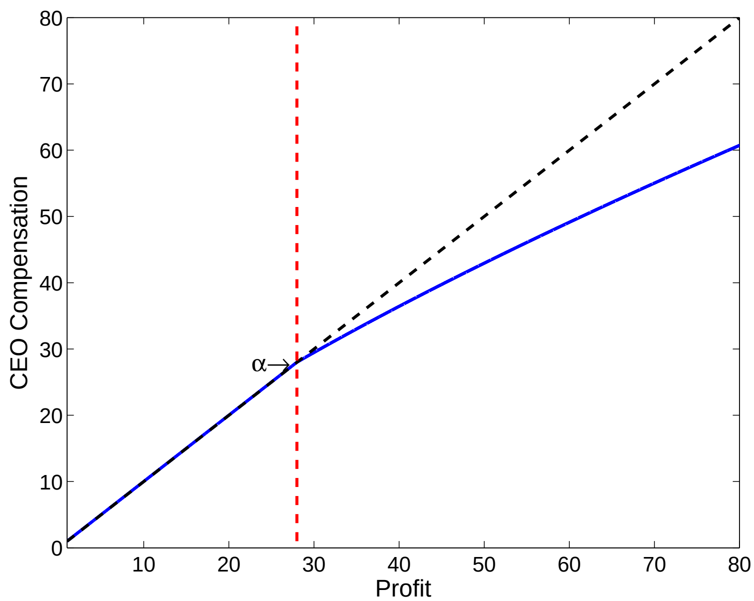

Figure 2.3: CEO Compensation Function

This graph plots the CEO compensation as a function of firm profit as defined in equation (2.10). $\alpha = 28.0$ and $\beta = 0.7373$. Firms with profit smaller than α are "sole proprietorship" firms while those with profit larger than α are "corporations."

I calibrate α to match the ratio of sales of all the corporations to the sales of all the firms, which is 62 percent in the U.S. in 2007.[24] $\beta \in (0, 1)$ is the elasticity of CEO income with respect to firm profit. To estimate this elasticity, I start with the ExecuCompustat using all the observations of U.S. public firms with non-missing values of compensation, stock ownership, and net income. I assume that the CEO's income is his/her total compensation plus a share of net income that is equal to his/her ownership of the firm and estimate the elasticity using this inferred income. β estimated with this approach is 0.73. This approach assumes that all the net income of a firm is distributed back to its shareholders, which is certainly an over-simplification. However, this approach provides a parsimonious way to capture the fact that the CEOs, both in the real world and the model world, are also significant owners of the firms they manage[25].

I impose an upper bound, s, on the human capital distribution to prevent the creation of unrealistically large corporations. I first compute the ratio between the highest CEO compensation in ExecuCompustat and the average U.S. wage from national income and product accounts (NIPA) in each year between 1992 and 2007[26]. I then calibrate $s = 2.8$ so that the same ratio in the model is matched to the median of the data sequence, which is around 2,903.

All the above parameters are fixed and reported in Table A.14. I jointly calibrate the iceberg trade cost and the fixed cost of starting foreign subsidiaries for each year between 1988 and 2008. I first assume that both cost matrices are sym-

[24] The sales of U.S. firms by legal form come from the *Statistics of U.S. Businesses, 2007* from the Census Bureau. The definition of "corporation" in this paper follows the legal form of "corporation" used by the Census. The other legal forms in the Census definition are classified as "proprietorship", which includes "S-corporations", "tax-exempt corporations", "partnership", "sole proprietorship", "other" and "tax-exempt other". The receipts of "government" are subtracted from the total firm sales.

[25] If I assume that CEO income equals CEO compensation, then the estimated β is around 0.36, which is within the range of traditional estimates of the size elasticity of CEO compensation (net of stock ownership returns). See *Frydman and Saks* (2010) and *Gabaix and Landier* (2008) for details of estimating the elasticity of CEO compensation (net of ownership returns) with respect to size.

[26] The wage data comes from NIPA Table 6.6A-D. The census does not allow disclosure of extreme values (maximum and minimum) that involve confidential data. Therefore I use the ratio between CEO compensation and the average U.S. wage instead of the CEO-to-worker pay ratio at the firm level in the empirical part.

metric: $\tau_{12} = \tau_{21}$ and $g_{12} = g_{21}$. Then I jointly calibrate the two costs $\{\tau_{21}, g_{21}\}$ to match the imports-to-GDP ratio and the multinational-firm-sales-to-GDP ratio in the U.S. in each year. The first moment condition can be directly estimated using the National Income and Product Accounts (NIPA). The data to estimate the second moment condition come from the Bureau of Economic Analysis's *Direct Investment and Multinational Corporations* data set.[27] These two parameters have to be jointly calibrated because iceberg trade costs affect not only the volume of trade but also the multinational sales through the extensive margin. Similarly, the fixed costs of FDI affect the volume of trade as well through the extensive margin. The calibrated sequence of trade barrier matrices is reported in Table A.16 in Appendix A.3. The calibration strategy creates a counterfactual world where only the volumes of trade and multinational sales match those observed in the real world, while all the other variables are roughly fixed between 1988 and 2008. In the next section, I examine how much income inequality the model is able to generate in this counterfactual world and how it compares to the real world.

2.6 Quantitative Results

2.6.1 Top Income Shares

Before turning into the distributional effect of globalization, I first present the model-generated top income shares and compare them with the U.S. data. In the benchmark model τ_{ij} and g_{ij} are calibrated to match the imports-to-GDP ratio and the multinational-sales-to-GDP ratio in each year between 1988 and 2008, while all the other parameters are fixed at values reported in Table A.14. Figure 2.4 compares the model-generated income shares with the data in 2008, and Figure A.1 presents

[27] I use "All non-bank foreign affiliates" sales data up to 2008 as the estimate for the sales of multinational firms.

the same comparison for the other years.[28]

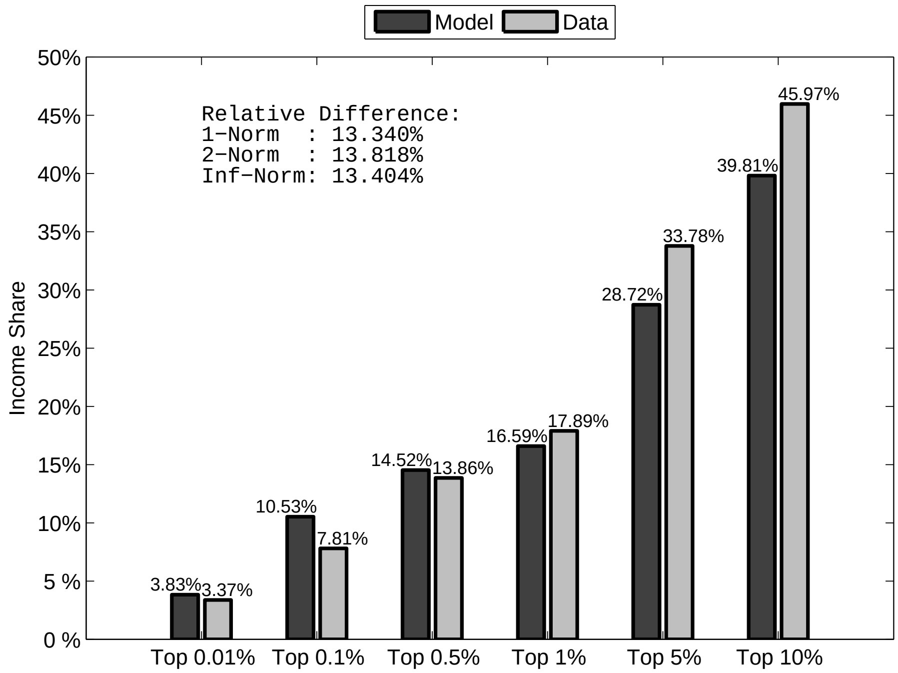

Figure 2.4: Top Income Shares: Model vs. Data (2008)

Note: This graph compares the top income shares between the model and the data in 2008. The top income shares in the model are described by the dark grey bars and those in the data described by light cray bars. The parameters behind the model simulation can be found in Section 2.5. The source of data is the updated Table A.1 from *Piketty and Saez* (2003). The average difference between the model and the data across the six top income shares is measured in Euclidean 1-norm, 2-norm, and infinity-norm. The differences are reported in percentage terms.

The model provides a good approximation of the U.S. income distribution, even though no parameter is calibrated to match any of the top income shares. For example, the top 0.01 percent income share is 3.37 percent in the data and 3.83 percent in the model in 2008. The top 1 percent income share is 13.86 percent in the data and 14.52 percent in the model. Overall, the difference between the model and the data for the six top income shares reported in Figure 2.4 is around 13 percent when measured in different Euclidean norms.

The model falls short in capturing certain income brackets. For example, the discrepancy in the top 0.1 percent income share is about 3 percentage points, and for

[28]The source of data is the updated Table A.1 from *Piketty and Saez* (2003).

the top 5 percent, the error accumulates to over 5 percentage points. In general, the explanatory power of the model declines when we move down the income ladder. This is because the model is not designed to explain the dynamics of income outside of the very rich. The key mechanism of the model is most suitable to explain the dynamics of income that is closely related to the performance of large firms. Outside of this group, the key mechanism is not directly applicable. For example, the individuals at the top 5 percent are usually highly paid working professionals (i.e., doctors, lawyers, engineers, and professors) in the real world. On the other hand, in the model, they are usually the CEOs at small domestic firms or workers with high human capital. It is important to understand how different professions are affected by the degree to which a country is exposed to the global markets, but this is well beyond the scope of this current project. The exclusive focus on executive compensation is also responsible for the model's inability to explain the response of income shares to globalization outside of the very rich, which is discussed later in this section.

2.6.2 Globalization and Firm-Level Inequality

I start the analysis of globalization on income inequality at the firm-level in general equilibrium. As in partial equilibrium, access to foreign markets widens the income gap between the CEO and the average workers.

To see this, we first compare the income of different individuals between autarky and trade. In "autarky," I set τ and g matrices arbitrarily high so that no trade and foreign investment takes place. In "trade," I calibrate the two matrices to match the data moments in 2008. The three panels in Figure 2.5 compare the income of the CEO and a worker with average human capital across three different firms in autarky and in trade. All three firms only sell to the domestic market in autarky. The firm in panel (a) remains a domestic firm in trade, the firm in panel (b) exports to the foreign market, and the firm in panel (c) serves the foreign market by FDI. The income of the

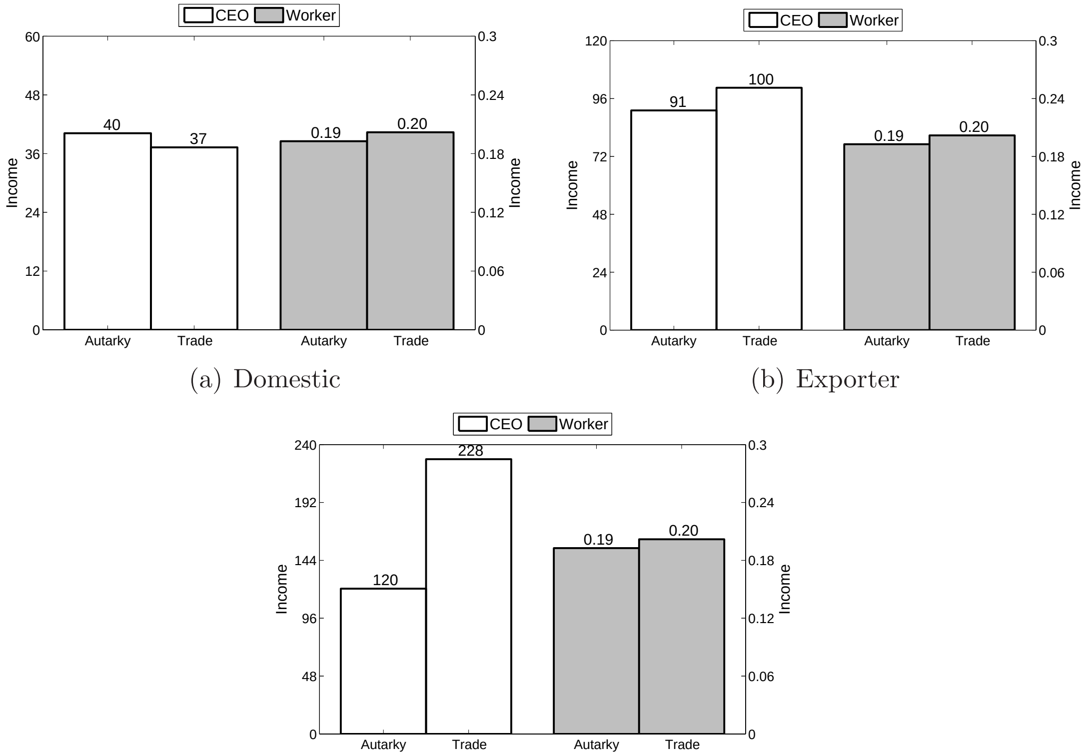

(a) Domestic

(b) Exporter

(c) MNE

Figure 2.5: CEO and Worker Income: Autarky vs. Trade

Note: This figure compares the income of the CEO and a worker with average endowment of human capital at three different firms in the economy. “Autarky” means both τ and g are set to a large number so trade and FDI fall to 0. “Trade” means both τ and g are calibrated so the imports-to-GDP ratio and multinational-sales-to-GDP ratio match the U.S. data in 2008.

average worker increases by around 5.3 percent in all three firms. The income of the CEO is tied to the performance of the specific firm, and therefore, the three different CEOs see different income paths. The CEO at the domestic firm sees his/her income decrease by around 7.5 percent, the CEO at the exporting firm sees his/her income increase by around 9.9 percent, while the CEO at the multinational firm sees his/her income surge by as much as 90 percent. Trade widens within-firm inequality for the firms that sell to the foreign market. In autarky, the CEO-to-worker pay ratio is 479 in the exporting firm, and it widens to 500 in trade. From autarky to trade, the CEO-to-worker pay ratio increases from 631 to 1,140 in the multinational firm.

Another way to visualize the gap-widening effect of trade is presented in Figure

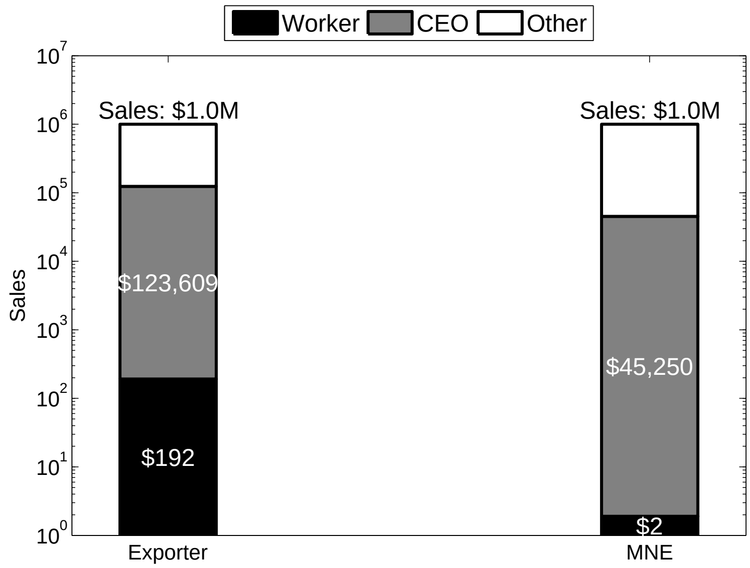

Figure 2.6: Revenue Distribution: Autarky v.s. Trade

Note: This figure shows how the income of the CEO and a worker with average human capital endowment changes if the firm's sales is increased by $1 million between autarky and trade for two different firms. "Autarky" means both τ and g are set to a large number so trade and FDI fall to 0. "Trade" means both τ and g are calibrated so the imports-to-GDP ratio and multinational-sales-to-GDP ratio match the U.S. data in 2008. The y-axis is in log-10 scale.

2.6. In this figure, I again compare the income profiles between autarky and trade in 2008. The two bars in the graph show how a $1 million increase in sales benefits the CEO and an average worker differently. For example, in a typical exporting firm (left bar), every $1 million increase in sales between trade and autarky is associated with a $123,609 increase in the compensation paid to the CEO and only a $192 increase in the compensation paid to an average worker. For a typical multinational firm (right bar), the distribution of "gains from trade" is more uneven. The CEO cuts $45,250 for himself/herself for every $1 million increase in sales, while the average worker earns only $2 more.

2.6.3 Globalization and Top Income Shares

How does firm-level inequality translate into economy-wide inequality? Figure 2.7 plots the income of the top 0.01 percent of the population in the model world in autarky and trade in 2008. The income distribution is concentrated to the right before the model economy opens up to trade, with the top 0.01 percent of the population claiming around 3 percent of total income. After opening up to trade, the distribution is even more skewed to the right, as the CEOs at the exporting and multinational firms cut a larger share from the extra profit earned abroad than the average workers. The surge in top income concentration can be observed as the gap between the red solid line (trade) and the blue dashed line (autarky) opens up. In the "trade" scenario, the top 0.01 percent income share increases to 3.83 percent. This is a 0.83 percentage point change in absolute income shares, or a 27.7 percent increase in relative terms. To put these numbers in perspective, the top 0.01 percent income share increased by 1.46 percentage points between 1970 and 1988 and another 1.38 percentage points between 1988 and 2008. Overall, the model seems to be able to explain a significant proportion of the change in top income share using the change in the volume of trade and FDI sales alone.

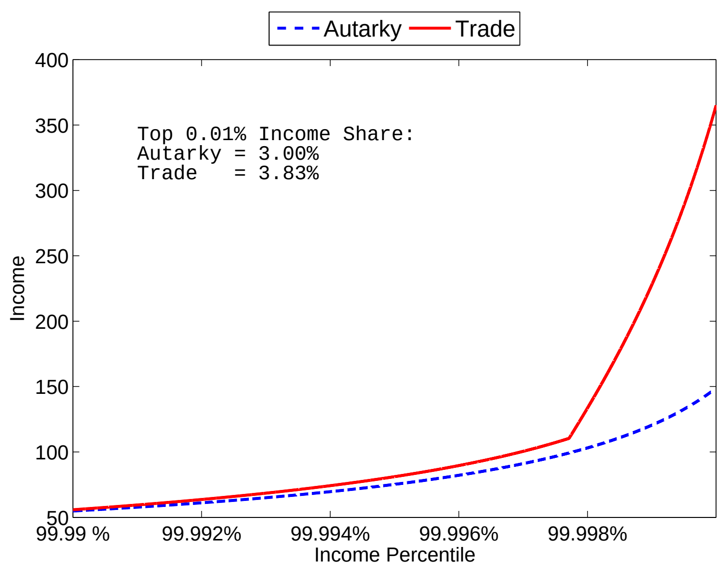

Figure 2.7: Top 0.01 percent Income Earners, Autarky vs. Trade

Note: This figure plots the income of top 0.01 percent in autarky v.s. in trade. "Autarky" means both τ and g are set to a large number so trade and FDI fall to zero. "Trade" means both τ and g are calibrated so the imports-to-GDP ratio and multinational-sales-to-GDP ratio match the U.S. data in 2008.

The calibration approach outlined in Section 2.5 creates a counterfactual world where only τ_{ij} and g_{ij} are allowed to move while all the other parameters are fixed. I will examine how much aggregate income inequality the model is able to generate in this counterfactual world and how it compares to the changes in the real world in the rest of this section.

The results of this exercise are summarized in Figure 2.8. This figure compares the change in the top 0.01 percent income shares between the model and the data between 1988 and 2008. The red dashed line (right axis) is the change of income shares between a given year in the x-axis and 1988 in the unit of percentage points in the data. For example, the last point on this curve indicates that comparing to 1988, the top 0.01 percent income share in 2008 is 1.38 percentage points higher. The blue solid line (left axis) is the same measure in the model in each year. I calibrate τ_{ij} and g_{ij} to match the imports-to-GDP ratio and multinational-sales-to-GDP ratio in each year, while keeping all the other parameters fixed as reported in Table A.14. Each point on the blue solid line is based on the top income share computed conditional on the calibrated τ_{ij} and g_{ij} in that year.

The model is able to capture the changes in top income shares over these 20 years. The correlation between the two curves in Figure 2.8 is 0.95, and the adjusted R-squared of regressing the data curve on the model curve is 0.89. In terms of magnitude, the changes in income shares in the model are on average a third of the data. For example, between 2008 and 1988 the top 0.01 percent income share increased by 1.38 percentage points in the data and 0.52 percentage points in the model, indicating that $0.52/1.38 \approx 37$ percent of the change in top income shares can be explained using the changes in trade volumes. The share of the data that can be explained by the model stays roughly the same and averages about 33 percent between 1998 and 2008. Many of the important factors that affect income inequalities are not included in the model mechanism at all, such as income tax incentives and equity markets. Nevertheless,

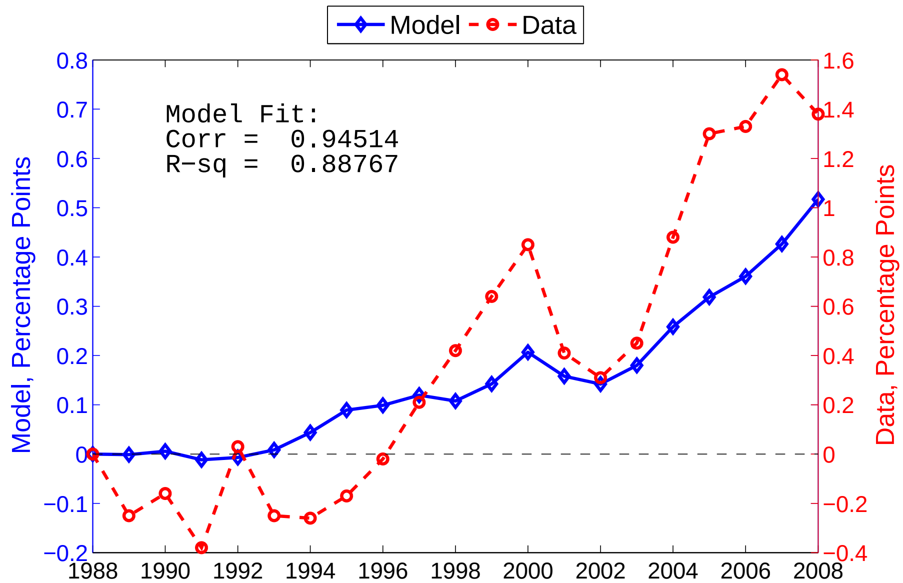

Figure 2.8: Income Share of the Top 0.01 Percent

Note: This graph shows the change in top 0.01 percent income shares in percentage points between 1988 and 2008. The change in the model is shown on the left axis and the change in the data is shown on the right axis. In the model simulation τ and g are calibrated to match the imports-to-GDP ratio and multinational-sales-to-GDP ratio in each year. For other model parameters behind this simulation, see Section 2.5. The source of the data is Table A.1 in the updated tables of *Piketty and Saez* (2003). Two measures of model fit are computed: the Pearson correlation between the two curves and the adjusted R-squared of estimating a linear relationship with data sequence on the left-hand-side and model sequence on the right (with constant term).

this counterfactual analysis indicates that a large proportion of the observed change in aggregate income inequality can be explained using the basic mechanism introduced by this paper: better market access introduced by globalization benefits the top executives and the average workers differently, widening the within-firm inequality.

Further reading of Figure 2.8 reveals more details. The changes in income shares in the data can be roughly separated into three phases. The first phase is from the beginning of the sample to around 1994, during which period the top income shares were volatile, mainly due to the short and long term effects of the 1986 Tax Reform Act (TRA).[29] This tax reform drastically changed the marginal tax rates and tax brackets for the top income earners, thus changing the tax reporting incentives significantly. The short-term consequences of the 1986 TRA are reflected in the sharp

[29] See *Slemrod* (1996) and *Poterba and Feenberg* (2000) for details.

increase in top income shares measured in the tax return data between 1986 and 1988 (not shown in the graph). The long-term consequences of the tax reform are less clear, but they can still be observed in the volatility of the data curve in Figure 2.8 before 1994. The model economy, in contrast, exhibits a steady increase in income shares, driven by trade and FDI sales. The ups and downs of the income shares in the data are not captured by the model because neither the income tax system nor the tax reporting incentives are modeled in this paper. In the second phase, starting from 1994, we start to observe a rapid increase in the top income share until the 2001-2002 stock market crash and economic recession. During this period, the surge in top income shares can be probably attributed to the rapid economic growth and the IT stock market boom. In the model world, the surge in income shares is less obvious. Again, the model is not designed to capture capital gains from the stock market and therefore misses the surge during this period. In the last phase from 2002 to the end of the period, we observe a strong surge in top income shares both in the data and in the model. This is a period during which globalization and inequality deepen at the same time: the imports-to-GDP ratio increases from 13 percent to 17 percent and the multinational-sales-to-GDP ratio increases from 27 percent to 47 percent in the data. During the same time, top 0.01 percent income share increased by roughly 1 percentage point, and it is captured by the model to a large extent.

The results of the same exercise for the top 0.1 percent of the income distribution are presented in Figure 2.9. The pattern is similar to that of the top 0.01 percent: the changes in income shares observed in the data are captured by the model over the entire period. The correlation between the two curves in the graph is 0.84 and the adjusted-R-squared is 0.68. On average, the share of the data that can be explained by the model is lower. Between 1988 and 2008, income share of the top 0.1 percent increased by 2.61 percentage points in the data, while it increased by 0.16 percentage points in the model. In other words, $0.16/2.61 \approx 6$ percent of the change in the data

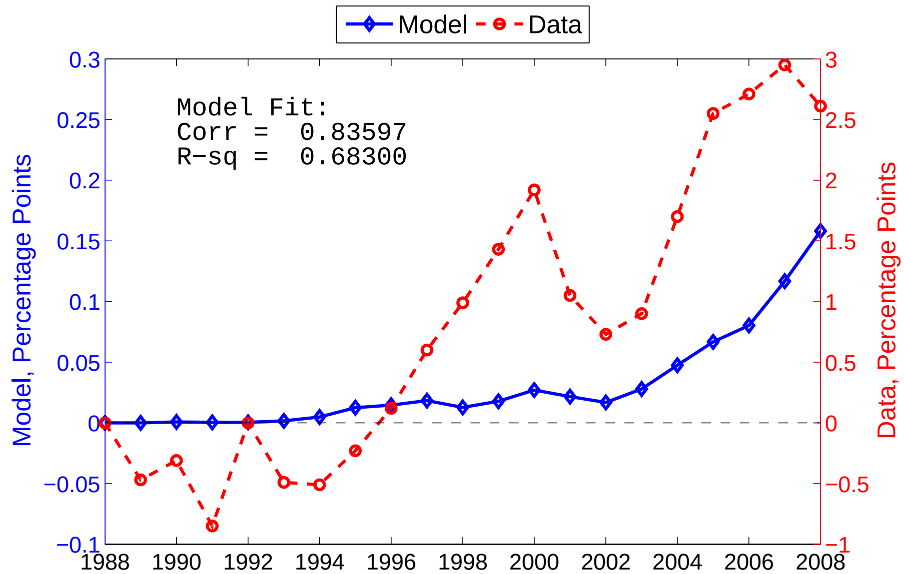

Figure 2.9: Income Share of the Top 0.1 Percent

Note: This graph shows the change in top 0.1 percent income shares in percentage points between 1988 and 2008. See the note to Figure 2.8 for more details.

can be explained, compared to 37 percent for the top 0.01 percent.

The approach outlined above makes a strong assumption that the changes in imports and multinational sales to GDP ratio are mainly determined by variable and fixed trade barriers. This assumption is based on the empirical findings that many key components of trade barriers have been declining over the past several decades, such as reported in *Hummels* (2007) and *Jacks et al.* (2008). The assumption is also motivated by the gravity models of international trade, where the variabilities in the volume of trade are driven by changes in trade costs. However, the main results of this paper do not depend on this assumption. Dynamics of income concentration at the top of the distribution depends on the size of the foreign market that firms can access to, but not on the exact channels through which the size of the foreign market is changed. Variable trade costs and fixed costs of starting foreign subsidiaries are parsimonious and empirically justified methods to simulate larger global markets, but by no means are they the only ways to do so.

The above counterfactual analysis shows that the expansion of trade volumes

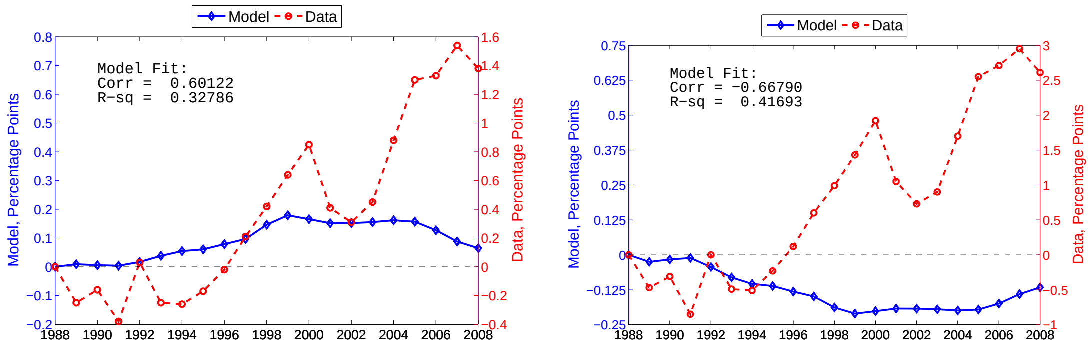

(a) Top 0.01 %

(b) Top 0.1 %

Figure 2.10: Top Income Shares, TFP Change

Note: This graph shows the change in top 0.01 percent and top 0.1 percent income shares in percentage points between 1988 and 2008. The change in the model is shown on the left axis, and the change in the data is shown on the right axis. In the model simulation, τ and g matrices are fixed at 1988 level, while TFP varies from year to year. For other model parameters behind this simulation, see Section 2.5. The source of the data is Table A.1 in the updated tables of *Piketty and Saez* (2003). Two measures of model fit is computed: the Pearson correlation between the two curves and the adjusted R-squared of estimating a linear relationship with data sequence on the left-hand-side and model sequence on the right (with constant term).

and multinational sales widens the income gap between the rich and the poor and drives up top income shares along the way. To understand the relationship between globalization and top income inequality further, I do another counterfactual analysis and show that without the expansion in trade and multinational sales, top income shares in the model will not exhibit the trends that we have observed in the data. In this counterfactual analysis I fix the τ and g matrices to the 1988 level in the previous example (the first row in Table A.16) and allow the estimated TFP vector b_i to vary. I compute b_i year by year using the methods outlined in Section 2.5. All the other parameters are fixed at values reported in Table A.14. Conditional on b_i, I solve the model and compute the top income shares for each year. In this counterfactual world, the volumes of trade and multinational sales barely move due to the fixed trade costs. The variations in top income shares are driven by the changes in TFP.

The results for top 0.01 percent and top 0.1 percent income shares are presented in two panels of Figures 2.10 in similar manners as in Figure 2.8 and 2.9. Without the expansion in the volume of trade and multinational sales, top income shares in the model do not follow the data. For example, the correlation between the data and model sequence is only 0.60 for the top 0.01 income share, compared to 0.95 in the case where trade moments are matched. The adjusted R-squared of regressing the data sequence on the model sequence is only 0.32, compared to 0.89 in the previous case. For the top 0.1%, the model generated income share even runs in the opposite direction with a correlation of -0.67 with the data.

2.7 Conclusion

This paper studies the relationship between globalization and income inequality with a special focus on the gap between the rich and the poor. Empirically, this paper presents a new fact that within-firm inequality is higher among the firms that have access to global markets. On average, the CEO-to-worker pay ratio is about 50 percent higher among the exporting firms than among domestic firms. The differences in within-firm inequality are mainly driven by differences in firm size. Exporting firms are more unequal because they are usually larger than their domestic counterparts.

This paper presents a new framework to study the distributional effect of trade. It merges the heterogeneous firms trade model with a model of occupational choice and executive compensation. The key mechanism to generate higher within-firm inequality among exporters and MNEs is through the size effect. On the one hand, CEO compensation is positively linked to the performance of the firm, and only the large and productive firms find it profitable to sell to the global markets. On the other hand, the wage rate is determined in a countrywide labor market and is not linked to each specific firm. These two forces imply that within-firm inequality is higher among the firms that have access to the global markets.

The paper argues that within-firm inequality can be responsible for a significant proportion of the surge of top income shares between 1988 and 2008. Using counterfactual analysis in which only the trade barriers are allowed to move exogenously, the model is able to broadly replicate the dynamics of top income shares. The correlation between the model-generated changes in top income share and the data is 0.95, and the adjusted R-squared is 0.89. In terms of magnitude, the changes in the model-generated income share are around 33 percent of the changes in the data. Similar but weaker results are found for the top 0.1 percent income share. These results suggest that globalization could have shaped the surge in top income shares in the U.S. through within-firm inequality significantly.

CHAPTER III

Lumpy Investment, Lumpy Inventories

Joint with Rüdiger Bachmann

3.1 Introduction

Researchers have now explored an ever more detailed and complex set of microeconomic frictions and heterogeneities in macroeconomic models. It has thus become an important question for macroeconomists, who on the one hand strive to build well-microfounded models, but are also, on the other hand, concerned about tractability and complexity of their models, how microeconomic frictions and heterogeneity affect macroeconomic dynamics. *Caplin and Spulber* (1987) present a striking example where any degree of nominal price stickiness at the micro level is consistent with the same aggregate outcome, money neutrality. In such a case, macroeconomic researchers arguably need not bother with the details of the microfoundation.

Conceptually, typical macroeconomic general equilibrium models can be split into a decision theoretic part where economic agents make often complex and dynamic decisions, which are, potentially, subject to a host of microeconomic frictions, e.g., physical adjustment frictions, informational frictions, etc. The second part of these models then consists of a formulation of aggregate resource and consistency constraints that will lead to the coordination of the individual decisions through prices (e.g., in Walrasian models) or aggregate quantities (e.g., in Non-Walrasian models, like search-and-matching models).

In this paper we argue that the answer to the question of how the microfoundations of decisions affect macroeconomic outcomes may depend on modeling choices in the second part, i.e., the details of how exactly general equilibrium closes a given physical environment, a perhaps obvious, but nevertheless underappreciated point. In other words, we will show – in a concrete, realistic and quantitative example – that there can be a cross effect between the general equilibrium part of a macroeconomic model and the mapping from microfoundations of decisions to macroeconomic outcomes.

Our example can simultaneously claim both realisticness with respect to a large body of microevidence (e.g., *Doms and Dunne* (1998) and *Cooper and Haltiwanger* (2006)) and also a certain notoriety in the literature: the debate about the aggregate importance of nonconvex capital adjustment costs. In a seminal paper, *Caballero and Engel* (1999) argue that nonconvex capital adjustment costs not only are powerful smoothers of aggregate investment, but also help explain certain nonlinearities in aggregate investment fluctuations. These results were produced in a macroeconomic model with essentially no general equilibrium elements, i.e., in a model with only a decision theoretic part that was aggregated by simple summation. In a series of papers, *Thomas* (2002), *Khan and Thomas* (2003) and *Khan and Thomas* (2008) argue, however, that once a general equilibrium part is added to the physical environment in *Caballero and Engel* (1999) not only do aggregate nonlinearities vanish, but also nonconvex capital adjustment costs have essentially no ability to smooth aggregate investment dynamics over and above what is done by general equilibrium price movements. Models with nonconvex capital adjustment costs thus deliver lumpy investment patterns at the micro level, but feature business cycle statistics that are identical to standard RBC models, once real wages and real interest rates adjust to clear markets.[1]

This somewhat striking irrelevance result can be understood from the first order conditions of the representative household, which are the same in a frictionless and a lumpy investment model, where the adjustment friction is on the firm side. With a representative household, the intratemporal and intertemporal first order conditions govern the optimal paths of consumption and labor supply, which in turn govern the optimal paths of output/income and saving in the short run. Thus, the households in a lumpy investment model *would like* to follow the same consumption path as in the frictionless model. The question is, whether they *are able* to do so when adjusting the capital stock is costly. The answer turns out to be yes, as long as the economy can substitute between the extensive and intensive margins of investment (see *Gourio and Kashyap* (2007) and, ultimately, *Caplin and Spulber* (1987) for this insight). To be concrete, after a positive aggregate productivity shock, the economy uses investment to increase consumption in the future. In a frictionless model this is entirely done through the intensive margin of investment: every firm invests a little more. With nonconvex capital adjustment costs this is no longer optimal, instead a few firms invest a lot. The desired amount of delayed consumption is concentrated into a few firms which really need to invest, and the same aggregate saving/investment path as in a frictionless model results.

This intuition rests on the assumption that the economy provides only one means of transferring consumption into the future, fixed capital. This is the familiar dual role of fixed capital in standard models: factor of production on the one hand and the only means of saving on the other, which in turn implies the familiar equality

[1] *Veracierto* (2002) makes a similar argument for kinked, but convex adjustment cost functions. *House* (2008) and *Miao and Wang* (2011) provide other sets of conditions on preferences, technology and the adjustment cost distribution under which fixed adjustment costs are neutral for business cycles dynamics. On the other side of the debate are *Gourio and Kashyap* (2007) and *Bachmann et al.* (2013), who argue that these irrelevance results are a matter of degree, specific to the calibration strategy used, and inconsistent with some nonlinear aspects of the time series of the aggregate investment rate in the U.S.

between saving and (fixed capital) investment. Thus for the economy as a whole investment and consumption dynamics are tightly linked. However, it is important to realize that this is only one particular way to introduce general equilibrium in a lumpy investment physical environment. There are others conceivable, and in reality an economy may delay consumption through multiple channels. We show that once we introduce multiple channels of investment and thus break the tight link between aggregate consumption and aggregate fixed capital investment, nonconvex adjustment costs and their magnitude matter much more for fixed capital investment dynamics in the sense that part of the partial equilibrium argument that they can act as smoothers of investment is restored. As has been mentioned above, this paper is very much about a cross-derivative from how the aggregate resource constraint is formulated to the ability of nonconvex adjustment costs to impact aggregate dynamics.[2]

The key intuition for this result is the substitution between different investment channels. Viewed from a social planners' perspective[3], introducing more investment channels offers more margins to smooth households' consumption, in addition to the extensive/intensive margin choice in fixed capital investment: if adjusting fixed capital is costly, the social planner can use other investment channels to optimally spread consumption over time. As a result, investment in fixed capital will be more sensitive to the level of frictions in capital adjustment.

To be concrete: we investigate the implications of multiple investment vehicles for the "neutrality question" in a quantitative DSGE model. Building on *Khan and Thomas* (2003) and *Khan and Thomas* (2007), we study a two-sector setting with an intermediate goods sector and a final goods sector. The final goods sector has the opportunity to store, at a cost, the output from the intermediate goods sector as

[2]To be clear, this is not a paper about aggregate investment nonlinearities, but rather about the ability of nonconvex adjustment costs to achieve what adjustment costs more generally are supposed to do: smooth, i.e., dampen and propagate, investment responses to shocks.

[3]We use in this paper a decentralized equilibrium model, where prices guarantee the social planners' optimal allocations, but for the intuition a social planners' perspective is useful.

inventories. The incentive to hold inventories is generated by fixed ordering costs for shipments from the intermediate goods to the final goods sector. The intermediate goods sector uses fixed capital as a production factor, whose adjustment is subject to nonconvex costs. We choose inventories as the second capital type because, 1) it is a highly cyclical component in the national accounts and, 2) it is a natural means to buffer consumption against temporary shocks. Methodologically, our paper provides the first quantitative analysis of how nonconvex capital adjustment frictions impact aggregate dynamics in the presence of capital good heterogeneity.[4]

Figure 3.1 summarizes the point of the paper in a nutshell. It shows the impulse response functions of fixed capital investment to a one standard deviation productivity shock. The nonconvex fixed capital adjustment costs dampen the initial response of fixed capital investment to a productivity shock by 2.99 percentage points in the presence of inventories ('Model I1' versus 'Model I2'). That is, the 'no capital adjustment costs'-impact response is approximately 50% higher than the one with capital adjustment costs. In contrast, without inventories nonconvex fixed capital adjustment costs dampen the initial response of fixed capital investment to a productivity shock by only 1.91 percentage points ('Model NI1' versus 'Model NI2'). That is, the 'no capital adjustment costs'-impact response is only 24% higher than the one with capital adjustment costs. This highlights the aforementioned interaction effect or cross-derivative, namely, that the presence of inventories, a second capital good, will quantitatively affect the difference between a frictionless and a frictional model for fixed capital. The difference here is measured in terms of one particular aggregate

[4] We push the research frontier in some dimensions, like the two-sector structure and inventories plus lumpy fixed capital investment. A paper related to ours is *Fiori* (2012), which also features lumpy capital adjustment in a two-sector model, but the focus there is on movements of the relative price of investment, which in our set up is constant by assumption. But we stay admittedly behind the research frontier in some other dimensions for reasons of computational tractability. For instance, unlike *Khan and Thomas* (2008) and *Bachmann et al.* (2013) we abstract from persistent idiosyncratic productivity shocks at the firm level. Micro heterogeneity is exclusively generated, just as in *Khan and Thomas* (2003), *Gourio and Kashyap* (2007) and *Khan and Thomas* (2007) by stochastic and ex-post different adjustment cost draws for both intermediate and final goods firms.

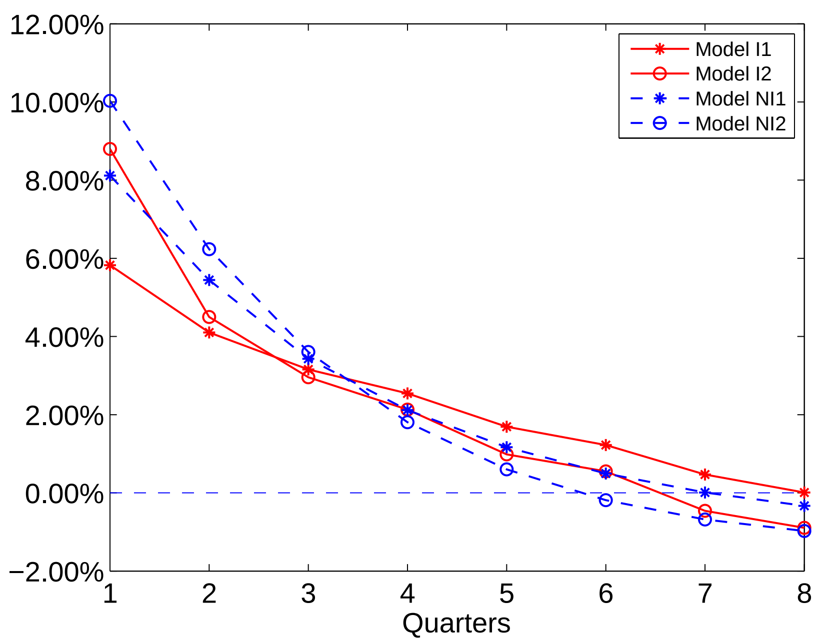

Notes: This figure shows the impulse response functions of fixed capital investment to a one standard deviation aggregate productivity shock in the intermediate goods sector. 'Model I1' has the baseline calibrated nonconvex fixed capital adjustment cost parameter and the baseline calibrated inventory order cost parameter. 'Model I2' has zero nonconvex fixed capital adjustment cost and the baseline inventory order cost parameter. 'Model NI1' has the baseline calibrated nonconvex fixed capital adjustment cost parameter and zero inventories. 'Model NI2' has zero nonconvex fixed capital adjustment cost and zero inventories. The difference between the IRFs of 'Model I2' and 'Model I1' is the effect of nonconvex fixed capital adjustment costs in the presence inventories. The difference between the IRFs of 'Model NI2' and 'Model NI1' is the effect of nonconvex fixed capital adjustment costs without inventories. There is no need to recalibrate the fixed capital adjustment cost parameter or the inventory order cost parameter, as our calibration targets, being long-run targets, are not sensitive across model specifications.

Figure 3.1: Impulse Response Function of Fixed Investment

statistic of interest: the initial response of fixed capital capital investment to an aggregate productivity shock. In addition, with inventories the response of investment in the model with the baseline level of nonconvex fixed capital adjustment costs is flatter than that in the model without these capital adjustment frictions. This means that with inventories nonconvex capital adjustment costs stretch the propagation of the productivity shock by more than what capital adjustment frictions can do without inventories. We will cast this argument in more quantitative terms below, when we study another important aggregate statistic and how its relation to nonconvex adjustment costs is shaped by the presence of inventories: autocorrelation coefficients of aggregate fixed capital investment.

Figure 3.1 also shows that inventories dampen the impact response of fixed capital investment at every level of fixed capital adjustment costs. With a positive productivity shock the higher demand for consumption transfer into the future can be partially satisfied by inventories, which are now relatively cheap to produce. And this is done the more so, the higher the nonconvex fixed capital adjustment is, i.e., the more costly the usage of fixed capital is: 10.01% impact response versus 9.01% impact response in the frictionless fixed capital adjustment model, yet 8.10% impact response versus 6.02% impact response in the model with the baseline calibrated nonconvex fixed capital adjustment cost parameter.

Another direct implication of our mechanism is, as we will show, that the households' ability to smooth consumption is enhanced when there are both inventories and fixed capital. In the end, inventories partially offset the hindering effect on consumption smoothing introduced by fixed capital adjustment frictions. As we will show, the impulse response functions of consumption to an aggregate productivity shock from the lumpy investment model and the frictionless adjustment model are very similar when inventories exist. Similarly, the volatility and persistence of aggregate consumption are much less sensitive to fixed capital adjustment frictions in models with inventories.

It is important to reiterate that the particular physical environment we chose – nonconvex capital adjustment costs as the friction and inventories as a way to modify the aggregate resource constraint[5] – are not as important as the general insight here: when aggregate resource constraints and general equilibrium effects are important for aggregate dynamics, the precise details of *how* these general equilibrium effects are introduced into the physical environment, the precise details of *how* the model is closed matter. In the words of *Caballero* (2010): "But instead, the current core approach of macroeconomics preserves many of the original convenience-assumptions from the research on the periphery[6] and then obsesses with closing the model by adding artificial factor supply constraints (note that the emphasis is on the word artificial, not on the word constraints)." This paper provides a quantitative analysis of the effects of closing the model in different ways for a specific, but prominent example. Put differently, unlike *Khan and Thomas* (2008) and *Bachmann et al.* (2013), who use the standard formulation for the aggregate resource constraint, this is not mainly a paper about the link between microfrictions and aggregate dynamics per se, but rather a paper about how this link is impacted by the formulation of the general equilibrium part of the model, i.e., a cross effect.

The rest of the paper proceeds as follows. Section 3.2 outlines the model. Section 3.3 discusses the calibration and model solution. Section 3.4 presents the results. Section 3.5 concludes.

[5]We conjecture that we could have used other ways of breaking the tight consumption-investment link in the standard model or used other functional form specifications for capital adjustment costs and gained similar insights.

[6]Caballero's terminology for the first, decision theoretic part of macro models.

3.2 The Model

3.2.1 The Environment

There are three kinds of agents in the economy: final goods producers, intermediate goods producers and households. The final goods producers use the intermediate goods, of which they hold inventories in equilibrium, and labor to produce the final goods.[7] Final output can be either consumed or invested as fixed capital. The intermediate goods producers combine fixed capital and labor to produce the intermediate goods. Households consume final goods and provide homogeneous labor to both types of producers. They own all the firms. They receive wage and dividend payments from both types of firms and purchase their consumption goods from the final goods producers. All markets are competitive.

3.2.1.1 The Final Goods Producers

There is a continuum of final goods producers. They use intermediate goods, m, and labor, n, to produce the final output through a production function $G(m, n)$. The production function is strictly concave and has decreasing returns to scale. Whenever the final goods producers purchase intermediate goods, they face a fixed cost of ordering and delivery, denoted in units of labor, ϵ. To avoid incurring the fixed cost frequently, the final good producers optimally hold a stock of inventories of the intermediate goods. Denote the inventory level for an individual producer as $s \in \mathbb{R}_+$.

The final goods producers differ in their fixed cost parameter for ordering, $\epsilon \in [0, \overline{\epsilon}]$. In each period, this parameter is drawn independently for every firm from a time invariant distribution $H(\epsilon)$. At the beginning of the period, a typical final firm starts

[7] To be clear on terminology: inventories in this model are not a capital good in the sense that they enter directly a production function, as in some modeling approaches in the literature. Thus, in our model, they lack the dual role of fixed capital. But they are a capital good in the sense that they represent a means of transferring consumption into the future, just like fixed capital. In this sense, we follow the NIPA terminology and denote net inventory changes as investment and the corresponding stock variables as capital.

with its stock of inventories, s, inherited from the previous period. It also learns its fixed cost parameter, ϵ. The firm decides whether to order intermediate goods. If the firm does so, it pays the fixed cost and chooses a new inventory level. Otherwise, the firm enters the production phase with the inherited intermediate goods inventory level s. We denote the quantity of adjustment by x_m. The inventory stock ready for production is $s_1 = s + x_m$, with $x_m = 0$ if the firm does not adjust.

After the inventory decision the firm determines its labor input, n, and the intermediate goods input, $m \in [0, s_1]$, for current production. Intermediate goods are used up in production. The remaining stock of intermediate goods, $s' = s_1 - m \geq 0$, is the starting stock of inventories for the next period. Stored inventories incur a unit cost of σ, denoted in units of final output. Inventory holding costs capture the idea that the storage technology that is used to partially circumvent the costly shipping technology is not free. Inventories require storage places, management and can lead to destruction of intermediate goods. The inventory management of the final good firms balances the trade-offs between costly shipping and costly storing optimally. In the end, the output of a typical final firm is $y = G(m, n) - \sigma s'$.

3.2.1.2 Intermediate Goods Producers

There is a continuum of intermediate goods producers. They are subject to an aggregate productivity shock, which is the sole source of aggregate uncertainty.[8] Let z denote the aggregate productivity level. It follows a Markov chain, $z \in \{z_1, \cdots, z_{N_z}\}$, where $P(z' = z_j | z = z_i) = \pi_{ij} \geq 0$ and $\sum_{j=1}^{N_z} \pi_{ij} = 1$ for all i.

Each firm produces with fixed capital and labor. Whenever the firm decides to adjust its capital stock, it has to pay a fixed cost, denoted in units of labor. In each

[8]As pointed out in *Khan and Thomas* (2007), placing aggregate productivity in the intermediate sector is necessary in this physical environment to generate a countercyclical relative price of intermediate goods, a feature found in the U.S. data. We abstract, for reasons of computability, from the persistent idiosyncratic productivity shocks that the recent literature has used to explain the observed micro-level heterogeneity in the data. Given the two firm problems, the computational burden in our model is already high, and we already push the computational frontier quite a bit.

period, the cost of adjusting capital is drawn independently for every firm from a time invariant distribution $I(\zeta)$. A typical intermediate good producer is identified by its capital stock, k, and its cost of adjusting capital, $\zeta \in [0, \overline{\zeta}]$.

At the beginning of each period, the firm learns aggregate productivity, z, and its idiosyncratic cost of adjusting capital, ζ. It starts with a fixed capital stock, k, inherited from the previous period. First, it decides about the labor input, l. It combines l and k according to a production function $zF(k, l)$. The $F(\cdot)$ function is strictly concave and has decreasing returns to scale.[9] After production, the firm chooses whether to adjust its capital stock. It can pay a fixed cost to adjust its capital stock by investing i. In this case, the new capital stock for the next period in efficiency units is $k' = [(1 - \delta)k + i]/\gamma$, where δ is the depreciation rate and γ is the steady state growth rate of the economy. Alternatively, the firm can avoid the adjustment cost and start the next period with the depreciated capital stock $k' = (1 - \delta)k/\gamma$.

3.2.1.3 Households

We assume a continuum of identical households who value consumption and leisure. They have access to a complete set of state-contingent claims. Households own all the firms. They provide labor to the firms and receive wage and dividend payments.

The households have the following felicity function:

$$u(c, n^h) = \log c - A^h n^h,$$

where n^h is the total hours devoted to market work.

[9] As *Miao and Wang* (2011) show, fixed adjustment costs cannot be expected to have a large impact with constant return to scale. We follow the majority of the literature, e.g., *Bachmann et al.* (2013), *Bloom* (2009), *Gourio and Kashyap* (2007) as well as *Cooper and Haltiwanger* (2006), and use a decreasing returns to scale assumption.

3.2.2 Competitive Equilibrium

3.2.2.1 Aggregate State Variables

In addition to z, the aggregate productivity level, two endogenously determined distributions are aggregate state variables in this model: the distribution of the firm-specific inventory stocks, $\mu(S)$, and the distribution of firm-specific fixed capital stocks, $\lambda(K)$. Both S and K are subsets of a Borel algebra over $\mathbb{R}_+$.

The aggregate state variables are summarized as (z, A), where $A = (\mu, \lambda)$. The distribution of μ evolves according to a law of motion $\mu' = \Gamma_\mu(z, A)$, and similarly, the distribution of λ evolves according to $\lambda' = \Gamma_\lambda(z, A)$.

The final good is the numeraire. Workers are paid $\omega(z, A)$ per unit of labor input. The intermediate goods are traded at $q(z, A)$ per unit.

3.2.2.2 Problem of the Household

The households receive a total dividend payment $D(z, A)$ and labor income $n^h(z, A)\omega(z, A)$ from the firms. In each period the households determine how much to work and how much to consume. All we need from the household problem is an intertemporal and an intratemporal first order condition.

We can express the dynamic programming problems for both types of firms with the marginal utility of consumption as the pricing kernel:

$$p(z, A) = \frac{1}{c(z, A)}.$$

Then every firm weighs its current profit by this pricing kernel and discounts its future expected earnings by β. This changes the unit of the firm's problems in both sectors to utils but leaves the policy functions unchanged.

The first-order conditions also imply that the real wage is given by:

$$\omega(z, A) = \frac{A^h}{p(z, A)}.$$

3.2.2.3 Problem of Final Goods Producers

Let V_0 be the value, in utils, of a final goods producer at the beginning of a period after the inventory adjustment cost parameter is realized and before any inventory adjustment and production decisions. Let V_1 be the expected value function after the adjustment decision but before the production decision. Given the aggregate laws of motion Γ_μ and Γ_λ, the firm's problem is characterized by the following three equations. For expositional ease, the arguments for functions other than the value functions are omitted.

$$V_0(s, \epsilon; z, A) = pqs + \max\Big\{ - p\omega\epsilon + V_a(z, A), -pqs + V_1(s; z, A)\Big\}, \tag{3.1}$$

$$V_a(z, A) = \max_{s_1 > 0}\{-pqs_1 + V_1(s_1; z, A)\}, \tag{3.2}$$

and:

$$\begin{aligned} V_1(s_1; z, A) = \max_{n \geq 0, s_1 \geq s' \geq 0} &\Big\{ p[G(s_1 - s', n) - \sigma s' - \omega n] \\ &+ \beta E_z \left[\int_0^{\overline{\epsilon}} V_0(s', \epsilon; z', A') d(H(\epsilon)) \right] \Big\}. \end{aligned} \tag{3.3}$$

The expectation is taken over z', next period's aggregate productivity.

Equation (3.1) describes the binary inventory adjustment decision of the firm. The firm adjusts if the value of entering the production phase with the optimally adjusted inventory level, described by $V_a(\cdot)$ in equation (3.2), minus the cost of adjustment, exceeds the value of directly entering the production phase with the inherited inventory level, $V_1(s; z, A)$.

The solution to equation (3.1) amounts to a cut-off rule in ϵ. The firm adjusts if:

$$-p\omega\epsilon + V_a(z, A) \geq -pqs + V_1(s; z, A).$$

Therefore the cut-off value is:

$$\widetilde{\epsilon}(s; z, A) = \frac{V_a(z, A) - V_1(s; z, A) + pqs}{p\omega}.$$

Given the support of the adjustment cost distribution, this cut-off value is modified to:

$$\epsilon^* = \max(0, \min(\overline{\epsilon}, \widetilde{\epsilon})).$$

The firm adjusts if its draw is smaller than or equal to $\epsilon^*(s; z, A)$.

Equation (3.2) describes the value of inventory adjustment. The solution to this equation is the optimal target level of inventory, $s_1^*(s, \epsilon; z, A)$. Note that the optimization problem, which is formulated in terms of the stock of inventories, s, instead of order flows, does not depend on any firm-specific characteristics. Therefore in any period, all the adjusting firms choose the same inventory target level, $s_1^*(z, A)$.

Equation (3.1) and (3.2) jointly determine the production-time inventory level, s_1:

$$s_1(s, \epsilon; z, A) = \begin{cases} s_1^*(z, A) & \text{if } \epsilon \leq \epsilon^*(s; z, A) \\ s & \text{if } \epsilon > \epsilon^*(s; z, A) \end{cases}.$$

Equation (4.3) describes the production phase. The firm finds the optimal inventory level for the next period and the optimal employment level for this period. The decision for next period's inventory level, s', is equivalent to deciding about the amount of intermediate goods to be used up in current production.

The solution for employment does not depend on the continuation value function. Therefore, given s', it is the analytical solution to:

$$\frac{\partial G(s_1 - s', n^*)}{\partial n} = \omega.$$

The optimal employment and inventory usage decision jointly imply the optimal output level:

$$y^*(s_1; z, A) = G(s_1 - s'^*(s_1; z, A), n^*(s_1; z, A)) - \sigma s'^*(s_1; z, A).$$

3.2.2.4 Problem of the Intermediate Goods Producers

Let W_0 be the value, in utils, of the intermediate good producers prior to the realization of the adjustment cost parameter ζ. Let W_1 be the value function after the realization of ζ. The intermediate good producer's problem can be summarized by the following equation:

$$\begin{aligned} W_1(k, \zeta; z, A) = \max_l \Bigg\{ & p \cdot [q \cdot zF(k, l) - l\omega] + \\ & \max \{W_i(k; z, A), -p\zeta\omega + W_a(k; z, A)\} \Bigg\}, \end{aligned} \tag{3.4}$$

where:

$$W_a(k; z, A) = \max_{k'} \{-(\gamma k' - (1-\delta)k)p + \beta E_z [W_0((k'; z', A')]\}, \tag{3.5}$$

$$W_i(k; z, A) = \beta E_z [W_0((1-\delta)k/\gamma; z', A')], \tag{3.6}$$

$$W_0(k; z, A) = \int_0^{\bar{\zeta}} W_1(k, \zeta; z, A) d(I(\zeta)). \tag{3.7}$$

The expectation in equation (3.5) and (3.6) is taken over z', next period's aggregate productivity.

In equation (3.4), the firm first solves for the optimal employment, given the fixed capital stock. The solution is:

$$\frac{\partial qzF(k,l^*)}{\partial l} = \omega.$$

After the production decision, the firm solves the binary fixed capital adjustment decision. The firm adjusts if the expected value from the optimally adjusted fixed capital stock, given in equation (3.5), minus the cost of adjustment, exceeds the expected value from the unadjusted fixed capital stock, given in equation (3.6).

The solution to the adjustment decision follows a cut-off rule for ζ. The firm adjusts if:

$$-p\omega\zeta + W_a(k;z,A) \geq W_i(k;z,A).$$

Therefore the cut-off value for ζ is:

$$\widetilde{\zeta}(k;z,A) = \frac{W_a(k;z,A) - W_i(k;z,A)}{p\omega}.$$

The restriction from the support of the cost distribution applies, so that

$$\zeta^* = \max(0, \min(\overline{\zeta}, \widetilde{\zeta})).$$

The firm adjusts to the target capital stock if its adjustment cost is smaller than or equal to $\zeta^*(k;z,A)$.

The optimal adjustment target for fixed capital is given by the solution to equation (3.5). Although the value function depends on the level of individual capital stocks, the resulting policy function, k^*, does not. After the binary adjustment decision, the capital stock for the next period is:

$$k'(k;z,A)=\begin{cases} k^*(z,A) & \text{if } \zeta\leq\zeta^*(k;z,A) \\ (1-\delta)k/\gamma & \text{if } \zeta>\zeta^*(k;z,A) \end{cases}.$$

3.2.2.5 Recursive Equilibrium

A recursive competitive equilibrium for the economy defined by:

$$\left\{u(c,n^h),\beta,F(k,l),G(m,n),\sigma,\delta,\gamma,H(\epsilon),I(\zeta),z\right\},$$

is a set of functions:

$$\left\{V_0,V_1,W_0,W_1,x_m,n,s',k',l,i,c,n^h,p,q,\omega,D,\Gamma_\mu,\Gamma_\lambda\right\},$$

such that:

1. Given ω, q, p, Γ_μ and Γ_λ, V_0 and V_1 solve the final firm's problem.
2. Given ω, q, p, Γ_μ and Γ_λ, W_0 and W_1 solve the intermediate firm's problem.
3. Given ω, D and p, c satisfies the household's first-order conditions.
4. The final goods market clears:

$$\begin{aligned} c(z,A)=&\int_S\int_0^{\overline{\epsilon}} y(s,\epsilon;z,A)d(H(\epsilon))d(\mu(s)) \\ &-\int_K\int_0^{\overline{\zeta}} i(k,\zeta;z,A)d(I(\zeta))d(\lambda(k)). \end{aligned}$$

5. The intermediate goods market clears:

$$\int_S \int_0^{\overline{\epsilon}} x_m(s,\epsilon;z,A)d(H(\epsilon))d(\mu(s)) =$$
$$\int_K \int_0^{\overline{\zeta}} zF(k,n(k,\zeta;z,A))d(I(\zeta))d(\lambda(k)).$$

6. The labor market clears:

$$n^h(z,A) = \int_S \int_0^{\overline{\epsilon}} \left(n(s;z,A) + \epsilon \cdot \mathbf{1}(x_m(s,\epsilon;z,A) \neq 0)\right) d(H(\epsilon))d(\mu(s))$$
$$+ \int_K \int_0^{\overline{\zeta}} (l(k,n(k;z,A)) + \zeta \cdot \mathbf{1}(i(k,\zeta;z,A) \neq 0))d(I(\zeta))d(\lambda(k)).$$

7. The laws of motion for aggregate state variables are consistent with individual decisions and the stochastic processes governing z:

   (a) $\Gamma_\mu(z,A)$ defined by $s'(s,\epsilon;z,A)$ and $H(\epsilon)$;

   (b) $\Gamma_\lambda(z,A)$ defined by $k'(k,\zeta;z,A)$ and $I(\zeta)$.

3.2.2.6 Some Terminology

Final Sales (FS), is defined as the total output of the final goods sector. Intermediate goods demand, X, is the total amount of intermediate goods purchased by the final goods sector. Intermediate goods usage, M, is the total amount of intermediate goods used up in production by the final goods sector. The difference between the two evaluated at the relative price of intermediate goods is Net Inventory Investment (NII):

$$\text{NII} = q \times (\text{X} - \text{M}).$$

Finally, Gross Domestic Product (GDP) in this physical environment is defined as the sum of final sales and net inventory investment:

$$\text{GDP} = \text{FS} + \text{NII}.$$

3.3 Calibration and Computation

3.3.1 Baseline Parameters

The model period is a quarter. We choose the following functional forms for the production functions:

$$F(k,l) = k^{\theta_k} l^{\theta_l},$$
$$G(m,n) = m^{\theta_m} n^{\theta_n}.$$

We discretize the productivity process z into $N_z = 11$ points following *Tauchen* (1986). The underlying continuous productivity process follows an AR(1) in logarithms with auto-correlation $\rho_z = 0.956$ and an innovation process with standard deviation $\sigma_z = 0.015$.

We set the subjective discount factor, $\beta = 0.984$, the depreciation rate $\delta = 0.017$, and the steady state growth factor $\gamma = 1.004$. A^h is calibrated so that the aggregate labor input equals 0.33. $\theta_m = 0.499$ is calibrated to match the share of intermediate inputs in final output. We set $\theta_k = 0.25$ and $\theta_l = 0.5$, the values used in *Bloom* (2009), which amounts to a capital elasticity of the firms' revenue function of 0.5[10]. We calibrate θ_n to match an aggregate labor share of 0.64. These parameters are summarized in Table 3.1:

[10] *Cooper and Haltiwanger* (2006), using LRD manufacturing data, estimate this parameter to be 0.592; *Hennessy and Whited* (2005), using Compustat data, find 0.551.

β	A^h	θ_m	θ_n	θ_k	θ_l	ρ_z	σ_z	δ	γ
0.984	2.128	0.499	0.367	0.250	0.500	0.956	0.015	0.017	1.004

Notes: β is the subjective discount factor of the households; A^h is the preference parameter for leisure; θ_m is the material share in the final good production function; θ_n is the labor share in the final good production function; θ_k is the capital share in the intermediate good production function; θ_l is the labor share in the intermediate good production function; ρ_z is the auto-correlation for the aggregate productivity process; σ_z is the standard deviation for aggregate productivity innovations; δ is the depreciation rate; γ is the steady state growth rate.

Table 3.1: Baseline Parameters

3.3.2 Inventory and Adjustment Cost Parameters

We assume that the inventory adjustment costs are uniformly distributed on $[0, \overline{\epsilon}]$. $\overline{\epsilon}$ is set so that the average inventory-to-sales ratio in the model equals 0.8185, the average of the real private non-farm inventory-to-sales ratio in the United States between 1960:1 and 2006:4. The unit cost of holding inventories, σ, is chosen so that the annual storage cost for all inventories is 12% of aggregate final output in value (see *Richardson* (1995) for details). These two targets jointly determine $\overline{\epsilon} = 0.3900$ and $\sigma = 0.0127$.

We assume that $I(\zeta)$ is uniform between $[0, \overline{\zeta}]$. The upper bound of the distribution is chosen so that the fraction of lumpy investors, defined as the firms whose gross investment rate is larger than 20% in a given year, is 18%. This calibration target is taken from *Cooper and Haltiwanger* (2006)'s analysis of manufacturing firms in the Longitudinal Research Database (LRD). This yields $\overline{\zeta} = 0.1841$.[11]

3.3.3 Numerical Solution

The inherent non-linearity of the model suggests global numerical solution methods. We use value function iterations from equation (3.1) to equation (4.3) to solve the problem of the final good producers. We use value function iterations from equa-

[11]It should be clear that the exact numbers for $\overline{\epsilon}$ and $\overline{\zeta}$ have little direct economic meaning and cannot be compared to other calibrations for these parameters in the literature. They are essentially free parameters to hit observable calibration targets (which are what is common across papers), such as the inventory-to-sales ratio and the fraction of firms that are lumpy investors. They will also lead to additional interpretable economic statistics like the average adjustment cost paid conditional on adjustment that we display below in Table 3.2. The precise values of these parameters are sensitive to the entire model environment and its calibration.

tion (3.4) to equation (3.7) to solve the intermediate firm's problem. Howard policy function accelerations are used to speed up convergence.

Our model gives rise to two endogenous distributions as state variables. We adopt the methods in *Krusell and Smith* (1997), *Krusell and Smith* (1998), *Khan and Thomas* (2003) as well as *Khan and Thomas* (2008) to compute the equilibrium. Denote the Ith moment of distribution $\mu(S)$ and $\lambda(K)$ as $\mu_I(S)$ and $\lambda_I(K)$ respectively. We approximate each distribution function with its first moment. We find that a log-linear form for the $\Gamma(\cdot)$ functions approximates the law of motion rather well in terms of forecasting accuracy:

$$\Gamma_\mu(z, \lambda_1, \mu_1) = \log \mu_1' = \alpha_\mu + \beta_\mu \log(\lambda_1) + \gamma_\mu \log(\mu_1) + \psi_\mu \log(z), \tag{3.8}$$

$$\Gamma_\lambda(z, \lambda_1, \mu_1) = \log \lambda_1' = \alpha_\lambda + \beta_\lambda \log(\lambda_1) + \gamma_\lambda \log(\mu_1) + \psi_\lambda \log(z). \tag{3.9}$$

We adopt similar rules for the pricing kernel and the relative price of intermediate goods:[12]

$$\log p = \alpha_p + \beta_p \log(\lambda_1) + \gamma_p \log(\mu_1) + \psi_p \log(z), \tag{3.10}$$

$$\log q = \alpha_q + \beta_q \log(\lambda_1) + \gamma_q \log(\mu_1) + \psi_q \log(z), \tag{3.11}$$

where λ_1 is the first moment of the capital stock distribution, and μ_1 is the first moment of the inventory stock distribution.

Given an initial guess for $\{\alpha_{\{.\}}, \beta_{\{.\}}, \gamma_{\{.\}}, \psi_{\{.\}}\}$, we solve the value functions as described above. Then we simulate the model without imposing the pricing rules in equations (3.10) and (3.11). In each model simulation period we search for a pair of prices, (p, q) such that all the firms optimize and all the markets clear under the

[12] We have experimented with other functional forms for the forecasting rules such as adding interaction terms between aggregate productivity and the capital and inventory moments. This did not lead to significant improvements in goodness-of-fit and often jeopardized numerical stability. Our specifications perform very well as measured by the R^2 of the equilibrium OLS regressions, which exceeds 0.9996 in all specifications.

forecasting rules in equation (3.8) and (3.9). To improve numerical accuracy, we use the value functions to re-solve all the optimization problems period by period and for every guess of (p, q). Given the market clearing prices, we update the capital and inventory stock distributions and proceed into the next period.

At the end of the simulation, we update the parameters $\{\alpha_{\{\cdot\}}, \beta_{\{\cdot\}}, \gamma_{\{\cdot\}}, \psi_{\{\cdot\}}\}$ using the simulated time series for the approximating moments and the market clearing prices. Then we repeat the algorithm with the updated parameters. Upon convergence of the parameters, we check the accuracy of the $\Gamma(\cdot)$ functions by the R^2 in the regression stage.

3.4 Results

We study the influence of nonconvex fixed capital adjustment costs on aggregate dynamics in our model by numerical simulation. We analyze four models that share all parameters other than $\overline{\epsilon}$ and $\overline{\zeta}$. 'Model I1' and 'Model I2' have the calibrated baseline equilibrium inventory holdings with $\overline{\epsilon} = 0.39$. 'Model I1' has calibrated fixed capital adjustment cost given by $\overline{\zeta} = 0.1841$, while 'Model I2' features a frictionless technology for adjusting the fixed capital stock. We also simulate two models without inventories, 'Model NI1' and 'Model NI2'. In these models, we set $\overline{\epsilon} = 0$ to eliminate equilibrium inventory holdings.[13] 'Model NI1' has the same level of $\overline{\zeta}$ as 'Model I1', while 'Model NI2' does not feature any frictions in adjusting the fixed capital stock. The parameter specifications for the four models are summarized in Table 3.2. We do not recalibrate $\overline{\zeta}$ in 'Model NI1' as the calibration targets are largely insensitive to the changes in equilibrium inventory levels, as shown in the fourth column of Table 3.2. To understand how the presence of inventories interacts with the effects of

[13]In theory, zero ordering costs are not inconsistent with positive inventory holdings as the firms might want to hedge against changes in the relative price of intermediate goods. However, in our simulations, given the inventory holding costs, no firm holds a positive level of inventories when $\overline{\epsilon} = 0$.

nonconvex fixed adjustment costs, we study the cross differences. That is, we contrast the differences between 'Model I1' and 'Model I2' with the differences between 'Model NI1' and 'Model NI2'.

Model Name	$\overline{\zeta}$	$\overline{\epsilon}$	Average Adjustment Cost	Fraction of Lumpy Adjusters	Note
I1	0.1841	0.3900	0.9300%	18.00%	Baseline fixed capital adjustment cost with inventory
I2	0.0000	0.3900	0.0000%	0.000%	Frictionless fixed capital adjustment with inventory
NI1	0.1841	0.0000	0.8900%	18.18%	Baseline fixed capital adjustment cost without inventory
NI2	0.0000	0.0000	0.0000%	0.000%	Frictionless fixed capital adjustment without inventory

Notes: 'Model I1' has the baseline calibrated nonconvex fixed capital adjustment cost parameter and the baseline calibrated inventory order cost parameter. 'Model I2' has zero nonconvex fixed capital adjustment cost and the baseline inventory order cost parameter. 'Model NI1' has the baseline calibrated nonconvex fixed capital adjustment cost parameter and zero inventories. 'Model NI2' has zero nonconvex fixed capital adjustment cost and zero inventories. "Average Adjustment Cost" is the average adjustment cost paid as a fraction of firms' output, conditional on adjustment. "Fraction of Lumpy Adjusters" is the share of lumpy adjusters, defined as the firms that adjust more than 20% of their initial capital stocks in a given year, in all firms.

Table 3.2: Model Specifications

We present four sets of results on those four models. We first compare their unconditional business cycle moments. Second, we study the impulse response functions for fixed capital investment and consumption across the four models. Third, we plot the volatility and persistence for consumption, fixed capital investment and, for the models with inventories, net inventory investment for a wider range of $\overline{\zeta}$. And finally, we analyze the role of general equilibrium price movements in bringing about these results.

3.4.1 Unconditional Business Cycle Analysis

After computing the equilibrium, we simulate the model for 1,000 periods, of which we discard the first 100 to eliminate the influence of initial conditions. Except for net inventory investment and fixed capital investment, all the simulated time series are transformed by natural logarithms and then detrended by an HP filter with smoothing parameter 1600. We detrend fixed capital investment with the HP filter directly and then divide the deviations by the trend. We divide net inventory investment by GDP and then apply the HP filter to this ratio.

The business cycle statistics in Panel (a) and (b) of Table 3.3 show several effects

(a) Standard Deviation

	GDP	Consumption	Fixed Investment	NII	Inventory Level
Model I1	1.4975	0.6416	9.6619	0.3793	1.2204
Model I2	1.5637	0.6336	11.5762	0.3240	1.1404
Model NI1	1.4772	0.7624	11.7371	-	-
Model NI2	1.5694	0.7436	13.8684	-	-
Data	1.6630	0.9015	4.8903	0.4220	1.6552

(b) First Order Auto-correlation

	GDP	Consumption	Fixed Investment	NII	Inventory Level
Model I1	0.6833	0.7623	0.7298	0.6157	0.9259
Model I2	0.6646	0.7932	0.6110	0.6616	0.9379
Model NI1	0.6839	0.7281	0.6648	-	-
Model NI2	0.6685	0.7739	0.6251	-	-
Data	0.8422	0.8833	0.9006	0.3696	0.8908

Notes: "NII" denotes net inventory investment. GDP, consumption, and inventory levels are logged and detrended with an HP filter with a penalty parameter of 1600. We detrend fixed investment with the HP filter and then divide the deviations by the trend. We divide NII by GDP and then detrend this ratio with the HP filter. All the standard deviations reported in Panel (a) are percentage points. Time period for the data moments: 1960:1 - 2006:4.

Table 3.3: Business Cycle Statistics

of inventories on aggregate dynamics.[14] The first message is that nonconvex fixed capital adjustment costs matter for aggregate dynamics. Business cycle dynamics differ significantly between 'Model I1' and 'Model I2'. For example, the percentage standard deviation of fixed capital investment decreases from 11.58 in the frictionless 'Model I2' to 9.66 in the lumpy investment 'Model I1'. Persistence of fixed capital investment increases from 0.61 to 0.73. In contrast, consumption volatility and persistence do not vary as much with the fixed capital adjustment cost parameter. Consumption dynamics are largely insulated from variations in capital adjustment frictions in the presence of inventories.[15]

Regarding the cross differences, the effects of nonconvex fixed capital adjustment

[14] *Bachmann et al.* (2013) is explicitly about how nonconvex fixed capital adjustment costs shape the implied model investment dynamics in terms of higher moments than standard second moments. They argue that aggregate investment data exhibits conditional heteroskedasticity and that micro nonconvexities are a natural mechanism to explain this. This paper is about micro nonconvexities and their role in shaping standard second moments and impulse response functions. Basically, this paper asks: is the fixed capital adjustment technology that is consistent with the micro data able to do what stand-in adjustment technologies do, namely, dampen and propagate aggregate investment.

[15] The excessively high fixed investment volatility, as shown in the third column of Panel (a), is a common property of two-sector models where fixed capital is only used in intermediate goods production. *Khan and Thomas* (2007) find similar results. As fixed adjustment cost works to dampen investment volatility, this might point to our calibration of $\overline{\zeta}$ being conservative, especially in light of the insights of *Bachmann et al.* (2013), who argue that focusing only on the fraction of lumpy investment episodes when calibrating nonconvex adjustment costs might lead to a downward biased estimate.

costs change significantly in models where inventories are absent. Most notably, the persistence of fixed investment only increases by 0.04 between 'Model NI2' and 'Model NI1', while it increases by 0.12 between 'Model I2' and 'Model I1'. The unconditional volatility of consumption increases by 0.0188 percentage points between 'Model NI2' and 'Model NI1' while it only increases by 0.0080 percentage points between 'Model I2' and 'Model I1'.[16] These results suggest that inventories strengthen the dampening and propagation effect of fixed adjustment costs on fixed capital investments.[17] At the same time, inventories enhance the households' ability to smooth consumption, making fixed capital adjustment costs much less effective in affecting consumption volatility.

As for net inventory investment and the level of inventories, we see that they behave exactly the opposite way from fixed capital investment, when the latter is subject to adjustment frictions. Their volatility rises and their persistence falls, when capital adjustment frictions are introduced. This is due to the substitution towards inventories as a means of consumption smoothing, as fixed capital becomes more costly to use.

3.4.2 Conditional Business Cycle Analysis - Impulse Response Functions

The first two panels of Figure 3.2 show the impulse response functions of aggregate fixed capital investment and consumption to a positive productivity shock in the intermediate goods sector. We simulate a shock process that starts with one standard

[16]Notice that the unconditional volatility of fixed investment drops roughly by the same amount in the cases with inventories and without. However, the unconditional volatility numbers are of course a combination of the increase in persistence and a decrease in conditional volatility, as shown in the impulse response function in Figure 3.1, and in the case of fixed investment both effects happen to roughly offset each other. In general, unconditional volatility can sometimes hide the interaction between conditional volatility and persistence, which is why we focus on the latter two in what follows.

[17]Note that already without inventories we have that nonconvex fixed capital adjustment costs matter somewhat for aggregate dynamics as, in line with the recent evidence in *Bloom* (2009) and *Cooper and Haltiwanger* (2006), our implied revenue elasticity of capital is closer to the calibration in *Gourio and Kashyap* (2007), where the substitution between the extensive and intensive margin of fixed capital investment is more difficult.

deviation above the median level of productivity, $z = 1$, and falls back to unity at the rate of $\rho_z = 0.956$.

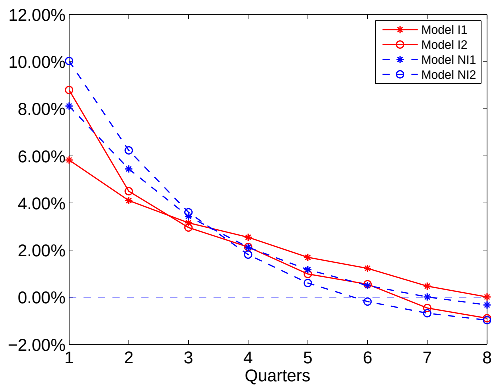

(a) Investment

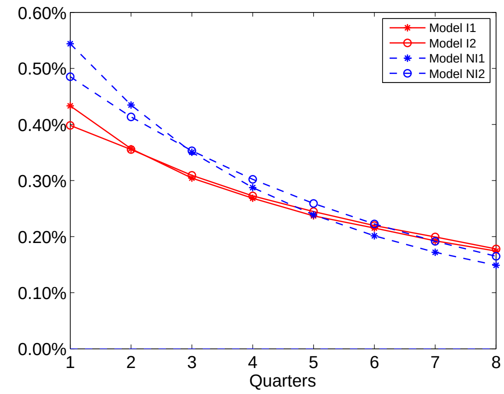

(b) Consumption

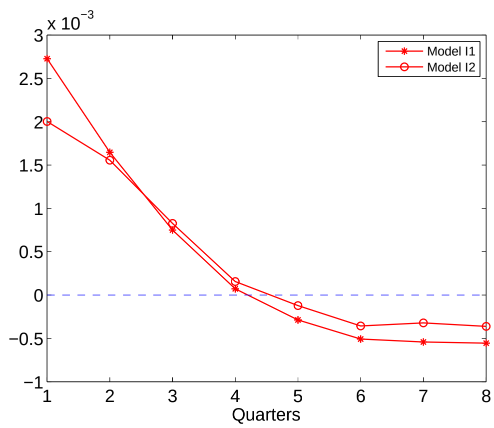

(c) Net Inventory Investment/GDP

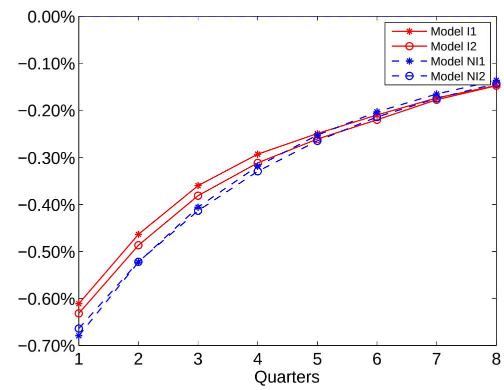

(d) Relative Price

Notes: This figure shows the impulse response functions of fixed capital investment, consumption, net inventory investment(NII) over GDP and the relative price to a one standard deviation aggregate productivity shock in the intermediate goods sector. 'Model I1' has the baseline calibrated nonconvex fixed capital adjustment cost parameter and the baseline calibrated inventory level. 'Model I2' has zero nonconvex fixed capital adjustment cost and the baseline calibrated inventory level. 'Model NI1' has the baseline calibrated nonconvex fixed capital adjustment cost parameter and zero inventories. 'Model NI2' has zero nonconvex fixed capital adjustment cost and zero inventories. The impulse response of net inventory investment over GDP is reported in absolute values, instead of percentage points, as the steady state value of net inventory investment is zero.

Figure 3.2: Impulse Response Functions

Fixed Capital Investment Panel(a) of Figure 3.2 presents the four impulse response functions for fixed capital investment. Comparing the models with $\overline{\zeta} = 0.1841$ against the models with $\overline{\zeta} = 0$ at the same level of inventories, we can see that nonconvex fixed capital adjustment costs dampen the initial responses both with and without inventories. However, at different levels of inventories, capital adjustment costs dampen these responses to a different degree. Without inventories, the initial response is dampened by 1.91 percentage points. In contrast, the initial response is dampened by 2.99 percentage points in models with inventories. Inventories also increase shock propagation. Comparing the impulse response function of 'Model I1' with that of 'Model NI1' without inventories, we see that the impulse response function in the model with inventories is flatter.

Both the extra dampening effect and the increased propagation of the shocks come from the key mechanism in our model: the substitution between fixed capital investment and inventory investment as a means of consumption smoothing. When adjusting fixed capital is costly, the economy switches to inventories. As a result, fixed capital investments do not need to respond to productivity shocks as much as when inventories are absent. The responses are also more protracted because firms tend to wait for lower adjustment cost draws to invest.

The flip side of the substitution between the two investment means can be observed in Panel(c) of Figure 3.2, which shows the impulse response functions of net inventory investment (over GDP). As expected, the response of net inventory investment is stronger when adjusting fixed capital investment is costly. In 'Model I1', the impact response is roughly 0.003, while in 'Model I2' it is only 0.002.[18]

The same mechanism can also explain the other cross effect, namely, how lumpy fixed capital investment changes the effect of inventories on aggregate investment dynamics. For both levels of fixed capital capital adjustment costs, inventories dampen the positive response of fixed capital investment to a positive productivity shock, as the latter is no longer used as much to ensure consumption smoothing. This switching away from fixed capital investment as a means of transferring consumption into the future is stronger, the more costly it is to use, i.e., when fixed capital adjustment

[18]The impulse responses for NII are reported in absolute changes as a fraction of GDP, not in percentage changes relative to the steady state. This is because the steady state value for NII is zero.

frictions are present. This explains why inventories dampen the initial response of fixed capital investment by somewhat over 2 percentage points with fixed capital adjustment frictions, but only by 1 percentage point, when fixed capital can be freely adjusted.

Consumption Another implication from the above mechanism is that consumers' ability to smooth consumption is enhanced by inventories. We illustrate this with the impulse response functions for consumption in Panel(b) of Figure 3.2.

First, the impact response from the models with inventories is below those from the models without inventories, for every level of fixed capital adjustment costs. Secondly, the smoothing effectiveness of inventories is so good that consumers despite the presence of capital adjustment costs can almost exactly recreate their frictionless consumption path. Nonconvex fixed capital adjustment costs barely change the response of consumption after the initial impact, when there are inventories. In contrast, without inventories nonconvex fixed capital adjustment costs do interfere with consumption smoothing.

We interpret these response functions as evidence that inventories provide an effective smoothing device for the consumers. As a result, consumption dynamics are less volatile when productivity shocks hit and capital adjustment frictions are less relevant for consumption dynamics in the presence of inventories.

3.4.3 Conditional Volatility and Persistence as a Function of Capital Adjustment Costs

In this section we illustrate the substitution mechanism between the two investment goods from a slightly different angle. We now simulate our model under our calibrated inventory level and the "No Inventory" setup over a wide range of $\overline{\zeta} \in [0, 0.4]$. The lower bound is frictionless adjustment, whereas the upper bound, 0.4, is approx-

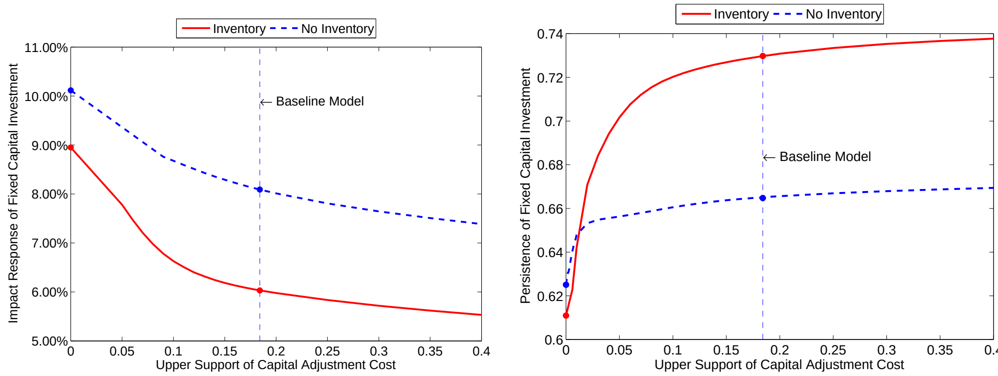

(a) Volatility (b) Persistence

Notes: This figure shows the impact response to an aggregate technology shock and the the first-order autocorrelation coefficient of fixed capital investment for models with $\overline{\zeta} \in [0, 0.4]$. The x-axis for both panels shows the upper bound of the capital adjustment cost distribution, $\overline{\zeta}$. In Panel(a), the y-axis shows the first element of the IRF of fixed capital investment to a one-standard deviation aggregate technology shock in percentage points. In Panel(b), the y-axis shows the first-order auto-correlation of fixed capital investment. For Panel(b) we detrend fixed capital investment with the HP(1600) filter and then divide the deviations by the trend.

Figure 3.3: Conditional Volatility and Persistence of Fixed Capital Investment

imately twice our baseline $\overline{\zeta} = 0.1841$.[19] We study how the conditional volatility, i.e., the impact response in the impulse-response function, and the persistence of fixed capital investment, consumption and net inventory investment change over this range of fixed capital adjustment costs.

Panel (a) of Figure 3.3 presents the conditional volatility of fixed capital investment over said $\overline{\zeta}$-range for both the inventory model and the "No Inventory" model. Independently of the level of inventories, higher capital adjustment costs dampen the impact response of fixed capital investment to aggregate shocks, and they do this in a more pronounced way in the model with inventories. The interaction between inventories and nonconvex capital adjustment costs is also apparent in the behavior of the persistence of fixed capital investment in Panel (b) of Figure 3.3. With inventories, persistence increases from 0.61 to 0.74 when $\overline{\zeta}$ changes from 0 to 0.4. In contrast, without inventories persistence only increases from 0.62 to 0.67 over the same range of $\overline{\zeta}$. The agents rely less on fixed capital investment when inventories are available.

[19] At $\overline{\zeta} = 0.4$ the annual fraction of firms which have lumpy investments is 15.23%, and the annual average adjustment cost paid conditional on adjustment and measured as a fraction of the firm's output is 1.66%.

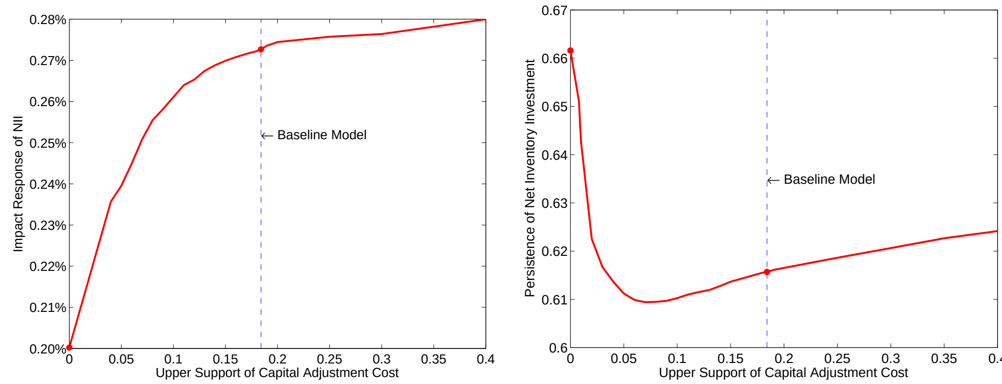

(a) Volatility (b) Persistence

Notes: See notes to Figure 3.3. This figure shows the impact response to an aggregate technology shock and the the first-order autocorrelation coefficient of net inventory investment (NII) divided by GDP for models with $\overline{\zeta} \in [0, 0.4]$.

Figure 3.4: Conditional Volatility and Persistence of Net Inventory Investment

As a result, the fluctuations in fixed capital investments are dampened and stretched. It is important to emphasize again that the central message of the paper lies in the different slopes of the two lines in both panels of Figure 3.3, which is precisely a graphical representation of the nontrivial cross effect between general equilibrium modeling and the impact of adjustment costs for fixed capital on aggregate statistics - conditional volatility and persistence.

We can directly observe the substitution between different investment channels by contrasting the conditional volatility of fixed capital investment in Figure 3.3 to the conditional volatility of net inventory investment in Panel (a) of Figure 3.4. As fixed adjustment costs increase, the agents rely more on inventories and less on fixed capital for consumption smoothing. As a result, higher fixed adjustment costs lead to more volatile net inventory investment and less volatile fixed capital investment. Panel (b) of Figure 3.4 shows the opposite, albeit with a small nonmonotonicity, effect on persistence of net inventory investment.

Also, we can see the implications of the investment substitution mechanism in the dynamics of consumption. Figure 3.5 shows that with inventories the conditional volatility of consumption is lower for every level of capital adjustment costs. More im-

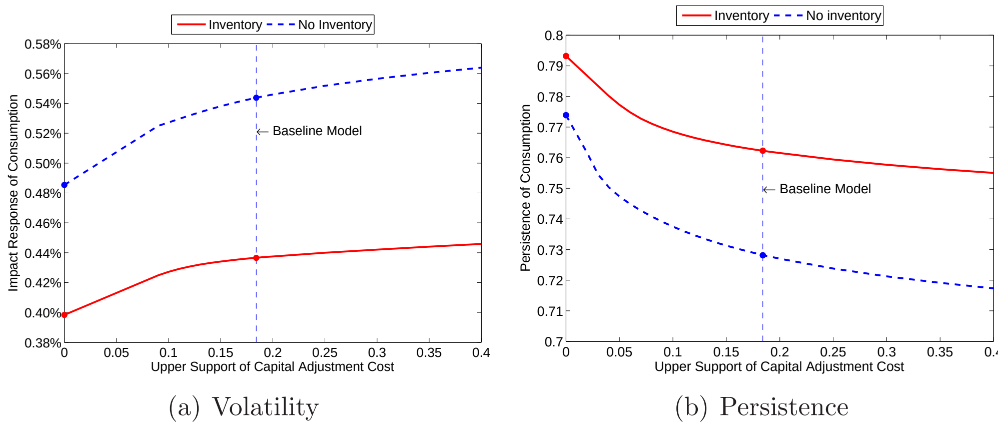

(a) Volatility

(b) Persistence

Notes: See notes to Figure 3.3. This figure shows the impact response to an aggregate technology shock and the the first-order autocorrelation coefficient of consumption for models with $\overline{\zeta} \in [0, 0.4]$. For Panel(b) consumption is logged and detrended with an HP filter with a smoothing parameter of 1600.

Figure 3.5: Conditional Volatility and Persistence of Consumption

portantly, as the slopes of the two curves suggest, the rate at which fixed adjustment costs increases conditional consumption volatility is lower when inventories exist. In other words, the same increase in fixed adjustment cost makes conditional consumption volatility move up higher when inventories are absent from the economy, whereas it can barely increase conditional consumption volatility when inventories are present.

The change in consumption persistence reveals the same mechanism, as shown in Panel (b) of Figure 3.5. The existence of inventories changes the degree to which fixed capital adjustment costs affect consumption persistence. Over the same range of $\overline{\zeta}$, consumption persistence decreases by much less in the inventory models compared to the "No Inventory" models.

3.4.4 The Effect of Market Clearing

The results on the effectiveness of fixed capital adjustment costs with or without inventories so far take into account all general equilibrium (GE) effects, i.e., adjustments of real interest rates and real wages, as well as the relative price of intermediate goods. In this section we isolate the effects of these price movements on how inventories impact the (ir)relevance of nonconvex fixed capital adjustment costs.

To this end, we solve three partial equilibrium versions of our model. In the first case, we fix both the pricing kernel, p, and the relative price q, at their long-run general equilibrium averages and simulate the model. In the second case, we fix the pricing kernel (and thus the real wage) to its long-run general equilibrium average, but allow the relative price to adjust so that the intermediate goods market clears. In the last case we fix the relative price to its long-run general equilibrium average, but allow the pricing kernel (and the real wage) to adjust so that the final goods market clears.

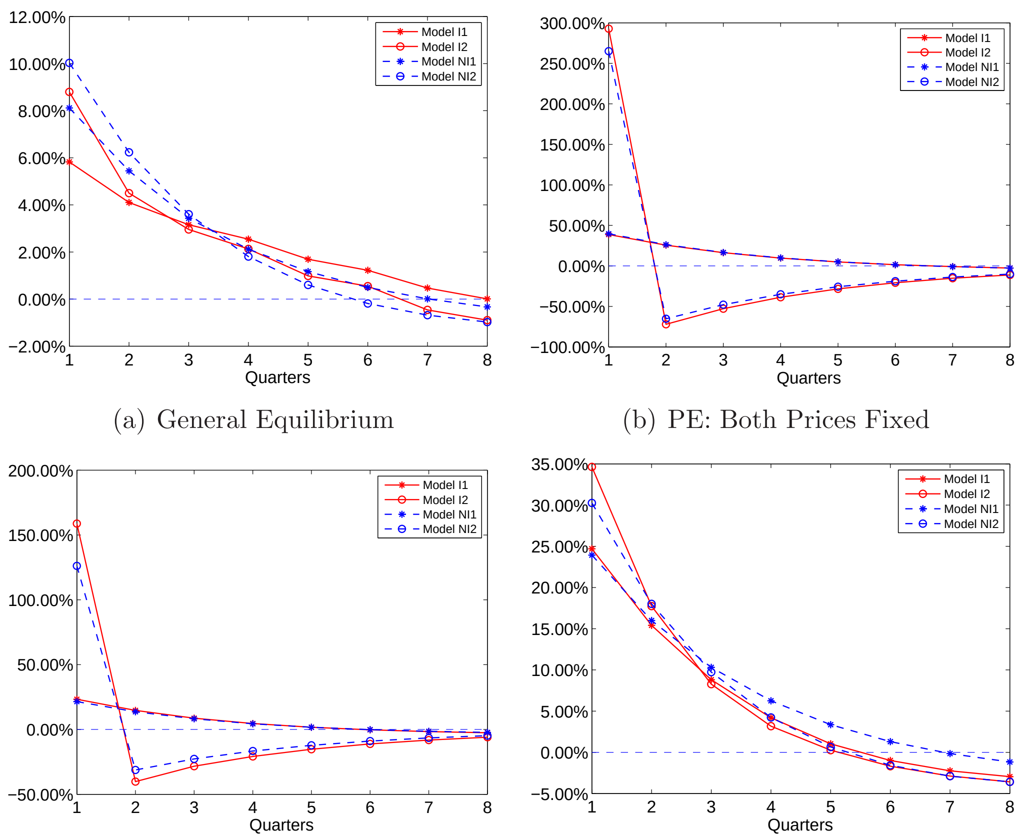

(a) General Equilibrium

(b) PE: Both Prices Fixed

(c) PE: Pricing Kernel Fixed

(d) PE: Relative Price Fixed

Notes: These are the impulse response functions for fixed capital investments. Panel(a) is the reproduction of Figure 3.1. Panel(b) is based on models where both the pricing kernel and the relative price are fixed. Panel(c) is based on models where only the pricing kernel is fixed. Panel(d) is based on models where only the relative price is fixed.

Figure 3.6: IRF for Fixed Capital Investments in Partial Equilibrium Models

The impulse response functions of fixed capital investment for all three cases are reported next to the full general equilibrium case – Panel (a) – in Figure 3.6. Panel(b) is the response from the first partial equilibrium case where both prices are fixed. Two messages emerge from this case. First, as is well known in the literature, nonconvex adjustment frictions matter a lot in partial equilibrium: the impact response drops substantively, and propagation arises only when fixed adjustment frictions are introduced. Second, inventories by and large do not change the effect of fixed adjustment frictions, as the differences between Model I1 and I2 are very similar to the differences between Model NI1 and NI2. Put differently, the effect of fixed capital adjustment frictions swamps the differential effect of inventories.

Panel(c) presents the response functions from the models where the pricing kernel is fixed but the relative price is not. The results in these models are very similar to those in the first case where both prices are fixed. Once again, nonconvex adjustment frictions matter a lot, but inventories do not interact with them significantly. Market clearing in the intermediate goods market only leads to slightly dampened fixed investment responses overall, as decreases in the relative price q (see Panel(d) of Figure 3.2) lead consumption smoothing activities away from fixed capital investment.

In other words, our exercise of comparing differences in differences really becomes only interesting, once real interest rate and real wage movements have been taken into account. The response functions in Panel(d) of Figure 3.6 come from the models where the pricing kernel and the real wage move freely to clear the final goods market, yet the relative price of intermediate goods is fixed. These response functions resemble those from the general equilibrium case in that in models with inventories the impact response of fixed investment is 40% higher with frictionless fixed capital adjustment, whereas in models without inventories it is only 26% higher.[20] Neverthe-

[20]The relative impact conditional on the same level of adjustment costs for fixed capital has changed between Panels (a) and (d) of Figure 3.6. For example, with no fixed capital adjustment costs, fixed capital reacts more to a productivity shock when there are inventories, but the price of intermediate goods is fixed, compared to the case where the price of intermediate goods adjusts downward, where the relative size of the reaction of fixed capital investment is reversed between the inventory and the 'no inventory'-case. Of course, with frictionless fixed capital adjustment, positive

less, market clearing in the intermediate goods market does play a role in rendering fixed capital adjustment frictions more relevant. Recall that in full general equilibrium the difference in the initial fixed investment response between the frictionless model and the lumpy model was 50% vs. 24%. The decline of the relative price q after an increase in aggregate productivity further facilitates the shifting of consumption smoothing through building up inventories and away from fixed capital investment. This substitution channel, for a given decline in q, is more valuable in an economy, when fixed capital adjustment is costly.

3.5 Conclusion

This paper shows that it matters for the aggregate implications of microfrictions *how* general equilibrium effects are introduced into the physical environment of dynamic stochastic general equilibrium models with these microfrictions. Specifically, we show that how relevant nonconvex fixed capital adjustment costs are for business cycle dynamics depends on how the aggregate resource constraint is modeled, depends on how the model is closed. Future research will explore the general insight in more general frameworks.

Here we develop a dynamic stochastic general equilibrium model to evaluate how the availability of multiple investment channels, here inventories in addition to fixed capital, affects the aggregate implications of nonconvex capital adjustment costs. We find that with more than one ways to invest, capital adjustment costs are more effective in dampening and propagating the response of fixed capital investment to an aggregate productivity shock.

inventory holding costs and a fixed price at which inventories can be stocked up, there is really not much reason to smooth consumption via inventories and thus fixed capital investment reacts more strongly. This changes, when the price of intermediate goods declines, which makes inventories as a smoothing device more attractive.

CHAPTER IV

Non-linearities in Aggregate Imports, Exports, and Real Exchange Rate Dynamics

4.1 Introduction

Aggregate U.S. imports, exports, and real exchange rates show conditional heteroscedasticity: they are more responsive to shocks when the past values of these variables are higher. Figure 4.1 presents a non-parametric estimate of the squared residual of the aggregate imports after fitting an autoregression model. It is clear from the figure that the standard deviation of the estimated residual increases with the lagged average aggregate imports.

In this paper, we document that the conditional heteroscedasticity of aggregate imports, exports, and real exchange rate is a pervasive feature of U.S. data. We estimate two ARCH family time series models and show that the non-linearity is statistically significant for imports and exports between 1970 and 2012, and for the real exchange rate between 1973 and 2012. Furthermore, we break down aggregate imports and exports into goods and services. We find that in most specifications, the imports and exports of goods exhibit conditional heteroscedasticity, while the trade in services does not.

These new empirical findings are important — they imply that the response of

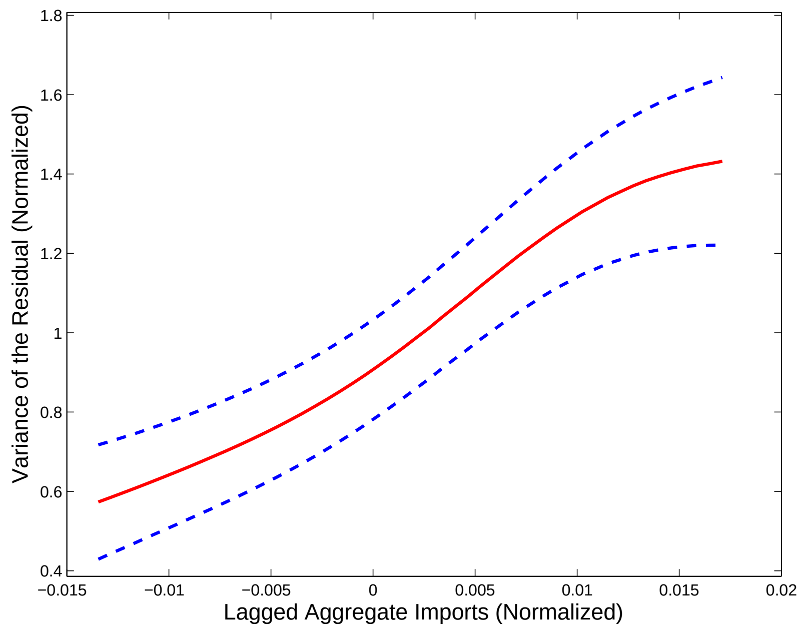

Figure 4.1: Conditional Heteroscedasticity of Aggregate Imports

Note: We use Gaussian kernels in the estimation. The variable on the x-axis is the lagged import-to-GDP ratio. The variable on the y-axis is the squared residual after fitting the data to an autoregression model. The lag in the autoregression model is 2, and the lag in the import-to-GDP ratio is 4. The choice of the lags is explained in Section 4.2. The dotted lines are one standard deviation bands. The y-axis is normalized so that the mean of the estimated function (solid red line) is 1. The x-axis is normalized so that the mean of the lagged average import-to-GDP ratio is 0. In this nonparametric estimation, the variance of the residual is 26 percent higher when the lagged aggregate imports are at the 75th percentile than at the median.

key international economics variables to policy interventions probably depends on the history of past shocks. Specifically, they suggest that the response of these variables might depend on whether we are at the peak or at the bottom of a business cycle. These empirical findings run against the predictions of many models in the international economics literature. For example, the international real business cycle (IRBC) model introduced in *Backus et al.* (1992) will predict a different pattern — the impact responses of imports, exports, and real exchange rates in an IRBC model only depend on the magnitude of the current shock, but not on the history of past shocks. At the end of the paper we discuss the possibility of incorporating inventory dynamics into a trade model, and the potential mechanisms that can be used to understand the new empirical findings.

Our work is related to the literature on conditional heteroscedasticity in macroeconomics. *Bachmann et al.* (2013) found that private fixed investment shows conditional heteroscedasticity in the United States. *Vavra and Berger* (2012) found similar patterns for aggregate durable expenditure. It is also broadly related to the literature on microeconomic lumpiness in investment and inventory, such as *Caballero and Engel* (1993), *Khan and Thomas* (2007), and *Bachmann and Ma* (2012).

The rest of the paper is organized as follows: Section 4.2 presents the time series models used to identify conditional heteroscedasticity. Section 4.3 provides the details on the data. Section 4.4 presents the main empirical findings. Section 4.5 concludes.

4.2 The Time Series Models

We use two time series models within the ARCH family to test whether U.S. imports, exports, and real exchange rate data exhibit conditional heteroscedasticity. The method is based on *Bachmann et al.* (2013). We first estimate the following auto-regression model with lag p:

$$x_t = \sum_{i=1}^{p} \phi_i x_{t-i} + \sigma_t e_t, \tag{4.1}$$

where x_t is the variable of interest, and e_t is an independently and identically distributed innovation term with a mean of zero and standard deviation of unity. σ_t is the key function to capture conditional heteroscedasticity. We model σ_t as a function of the kth lag moving average of x_t, denoted as $\overline{x}_t^k$:

$$\overline{x}_t^k = \frac{1}{k} \sum_{i=1}^{k} x_{t-i}.$$

We propose two functional forms for σ_t. For Model 1, we use the following:

$$\sigma_t = \alpha_1 + \eta_1 \bar{x}^k_{t-1},$$

and for Model 2, we use the following:

$$\sigma_t^2 = \alpha_2 + \eta_2 \bar{x}^k_{t-1}.$$

The test of conditional heteroscedasticity is equivalent to the following test:

$$\mathrm{H}_0 : \eta_1 = \eta_2 = 0.$$

Rejection of the null hypothesis implies that σ_t will comove with $\bar{x}^k_t$. Conversely, if η_1 and η_2 are not significantly different from zero, both models are reduced to standard autoregressions with homoscedasticity.

Conditional heteroscedasticity is closely related to the history-dependency of impulse response functions (IRF) in this context. Equation (4.1) implies that the impulse response of x to e upon impact at time t, denoted by $\mathrm{IRF}_{0,t}$, is σ_t. Therefore, we can express the impact response as follows:

$$\mathrm{IRF}_{0,t} = \begin{cases} \alpha_1 + \eta_1 \bar{x}^k_{t-1} & \text{for Model 1,} \\ \\ \sqrt{\alpha_2 + \eta_2 \bar{x}^k_{t-1}} & \text{for Model 2.} \end{cases}$$

If η_1 or η_2 are significantly different from zero, the above equation implies that the IRF of x depends on the past values of x, and subsequently, depends on the past history of shocks as well. Conversely, if η_1 and η_2 are not significantly different from zero, then the IRF depends only on the size of the current shock, characterized by α_1 and α_2.

In practice we use a two-step procedure to estimate the above model. Assume that we have the quarterly import-to-GDP ratio, x_t, with $t = 1, 2, \cdots, T$. For a given combination of lags $\{p, k\}$, we first estimate equation (4.1) using OLS. We then use the residuals, ϵ_t, from the OLS to estimate η with the following specification for Model 1:

$$|\epsilon_t| = \sqrt{\frac{2}{\pi}}(\alpha_1 + \eta_1 \bar{x}^k_{t-1}) + u_t, \tag{4.2}$$

and the following specification for Model 2:

$$\epsilon_t^2 = \alpha_2 + \eta_2 \bar{x}^k_{t-t} + u_t. \tag{4.3}$$

We determine the optimal lags of p and k using the Bayesian Information Criterion (BIC). To this end, we first set a p_{max} and a k_{max} and then estimate the models for all the combinations of $\{p, k\}$ such that $p \leq p_{max}$ and $k \leq k_{max}$. We then choose the combination of $\{p, k\}$ with the lowest BIC. To ensure that the models are comparable, we use the same sample size for all the estimations, so that the number of observations equals to $T - \max\{p_{max}, k_{max}\}$ for all combinations of $\{p, k\}$. We pick p_{max} and k_{max} to ensure that the optimal p and k are strictly smaller than these upper bounds. Unless otherwise noted, we set $p_{max} = k_{max} = 15$.

We measure conditional heteroscedasticity with several statistics. The significance of the estimated η is measured with the Student's t statistics. We also measure the magnitude of the non-linearity using the spread of the predicted order statistics. For a given pair of estimated $\{\alpha, \eta\}$, we predict the values of ϵ using equation (4.2) and (4.3) for Model 1 and 2 respectively. We then measure the spread of the predicted ϵ using the ratios between different percentiles of the predicted ϵ. Specifically, we report the ratio between the 95th and the 5th percentile, and the ratio between the 90th and the 10th percentile.

We bootstrap the model to estimate the p-values associated with the t statistics. The bootstrapping algorithm works as follows. Given a set of estimated $\{\alpha, \eta\}$, we first predict ϵ using equation (4.2) and (4.3) for Model 1 and 2 respectively. Denote the predicted sequence of ϵ_t as $\hat{\epsilon}_t$. We then back-out the i.i.d. innovations e_t as follows:

$$e_t = \frac{\epsilon_t}{\hat{\epsilon}_t},$$

where ϵ_t is the residual from the OLS estimation of equation (4.1) at the first stage. We bootstrap on the sequence of e_t. Based on the bootstrapped e_t, we generate a new sequence of x_t using equation (4.1) and the estimated ϕ. The length of the new x_t is $T + D$, where T is the original length of the x_t sequence. The first D observations of the bootstrapped x_t are discarded to ensure that the new sequence of x_t does not depend on the initial values. At last the entire model is estimated again with the bootstrapped x_t, and a new pair of $\{\alpha, \eta\}$ is generated. Given a number of bootstrapped η, we estimate the p-value of the t statistics as usual. In practice we bootstrap the model 200 times and set $D = 100$.

4.3 Data

The aggregate imports and exports data come from the National Income and Product Accounts (NIPA).[1] The variables of interest are the import-to-GDP ratio and the export-to-GDP ratio at the quarterly frequency. We use the data from the 1st quarter of 1970 to the 4th quarter of 2012. Over this period, both the import-to-GDP ratio and the export-to-GDP ratio show a strong upward trend. We apply a linear filter to these two sequences to ensure that the data series are stationary when estimated with the autoregression models. We also use the Hodrick-Prescott (HP)

[1] Table 1.1.5, seasonally adjusted gross domestic product and its components at annual rates.

filter with a weight of 1600 as robustness check.

The exchange rate data come from the Board of Governors of the Federal Reserve System. We use the real trade-weighted U.S. dollar index for major currencies at the quarterly frequency.[2] The index is available from the 1st quarter of 1973 to the 4th quarter of 2012. The exchange rate index does not show a strong time trend over the period; therefore, we report the results based on unfiltered data. We also report the results based on linearly-filtered data as robustness check.

4.4 Results

4.4.1 Imports

Table 4.1 presents the results based on a linear filter for U.S. aggregate imports from the 1st quarter of 1970 to the 4th quarter of 2006. The first two rows of Table 4.1 report the optimal p^* and k^* chosen by BIC. The third and fourth rows report the estimated η and the associated Student's t statistics. The associated p values for the t statistics estimated with bootstrapping are reported in the fifth row. The last two rows report the magnitude of non-linearity using the ratios between different percentiles of the predicted $\hat{\epsilon}$. $\hat{\epsilon}_m$ denotes the mth percentile of the predicted residual.

The first two columns of Table 4.1 show that U.S. aggregate imports exhibit conditional heteroscedasticity. The estimated η for both models are significantly positive, measured either by the t statistics or the bootstrapped p value. The magnitude of heteroscedasticity is large. For example, the 95th percentile $\hat{\epsilon}$ is 180 percent higher than the 5th percentile ($e^{1.032} \approx 2.80$) in Model 1; the 95th percentile $\hat{\epsilon}$ is 179 percent higher than the 5th percentile ($e^{1.025} \approx 2.79$) in Model 2.

We break down the aggregate imports into goods and services and repeat our analysis. The variables of interest here are the imports of goods and services, again

[2]Currencies included in the index are: Euro, Canadian dollar, Japanese yen, British pound, Swiss franc, Australian dollar, and Swedish krona.

Model	All 1	All 2	Goods 1	Goods 2	Services 1	Services 2
p^*	2	2	2	2	1	1
k^*	4	4	5	5	8	8
$\eta \times 1000$	112.6	0.5614	123.2	0.5567	64.48	0.04121
t	4.418	4.029	4.757	4.131	1.527	1.274
p-value-bootstrap	0.0002124	0.005435	0.0005399	0.004786	0.1007	0.1534
$\pm\log(\epsilon_{95}/\epsilon_5)$	1.032	1.025	1.194	1.15	0.363	0.2553
$\pm\log(\epsilon_{90}/\epsilon_{10})$	0.795	0.7235	0.8929	0.7671	0.3183	0.2241
No.Obs	133	133	133	133	133	133

Table 4.1: U.S. Aggregate Imports, 1970 – 2006, Linear Filter

Note: This table reports the estimation of equation (4.1), (4.2), and (4.3) using U.S. aggregate imports data. Data source: NIPA Table 1.1.5. Time frame is 1st quarter of 1970 to the 4th quarter of 2006. The variables of interest are linearly-filtered import-to-GDP ratios.

normalized by GDP. The next four columns of Table 4.1 report the results. The non-linearity found in aggregate imports is mainly driven by the imports of goods, but not services. The estimated η of goods is at the same level of magnitude as that of the aggregate imports, while the estimated η for services is significantly smaller. For example, in Model 1, the estimated η is 123.2 for the imports of goods, 112.6 for aggregate imports, while it is only 64.48 for the imports of services. Moreover, the estimated η for the imports of services is not significantly different from zero. The bootstrapped p-value for Model 1 is at 0.10 and at around 0.15 for Model 2. This suggests that, different from the imports of goods, the imports of services seems to be conditional homoscedastic.

The data used in Table 4.1 stop at 2006 before the Great Recession. The trade collapse following the Great Recession introduces a high level of volatility to the data, which changes the estimation of η significantly. In practice, this creates problems for our bootstrapping algorithm and the predicted order statistics. The root of the problem is as follows. The left hand sides of equations (4.2) and (4.3) are, by definition, always non-negative. However, the right hand sides of these two equations are linear, and thus can potentially take negative values. In most of the estimations of $\{\alpha, \eta\}$,

Model	All 1	All 2	Goods 1	Goods 2	Services 1	Services 2
p^*	1	1	1	1	1	1
k^*	5	1	3	1	5	5
$\eta \times 1000$	168.4	1.151	134.9	0.9495	33.41	0.0291
t	4.695	2.128	3.888	1.764	1.334	1.383
p-value-bootstrap	0.0006336	$\underline{0.06548}$	0.02914	$\underline{0.1163}$	0.1991	$\underline{0.1681}$
$\pm\log(\epsilon_{95}/\epsilon_5)$	1.541	$\underline{1.391}$	1.206	$\underline{1.003}$	0.2925	$\underline{0.2735}$
$\pm\log(\epsilon_{90}/\epsilon_{10})$	1.185	$\underline{1.166}$	0.9992	$\underline{0.8487}$	0.2611	$\underline{0.2448}$
No.Obs	157	157	157	157	157	157

Table 4.2: U.S. Aggregate Imports, 1970 – 2012, Linear Filter

Note: This table reports the estimation of equation (4.1) and (4.2) using U.S. aggregate imports data. Data source: NIPA Table 1.1.5. Time frame is 1st quarter of 1970 to the 4th quarter of 2012. The variables of interest are linearly-filtered import-to-GDP ratios. See the main text for the explanation of the statistics with underlines.

this will not happen. However, when the data after 2006 are included, we tend to get more extreme estimations of α and η, and subsequently, negative predictions of $\hat{\epsilon}^2$. This does not affect the point estimate and the t statistics, but it does affect the estimated spread of $\hat{\epsilon}$ and the bootstrapped p-value. We work around this problem by forcing a corner solution, and assume $\hat{\epsilon}_t = e_t = 0$ whenever the predicted $\hat{\epsilon}^2 \leq 0$. As this is essentially a different bootstrapping strategy, the comparison between these results and the previous ones should be made with caution. We underline the statistics derived using the above-mentioned work-around in Table 4.2. Note that in practice we do not get negative predictions of $|\epsilon_t|$ in Model 1, and therefore the estimations of p-value and the order statistics are comparable to the previous table.

The qualitative results with data from 1970 to 2012 are similar to those that stop at 2006: both the aggregate and the goods imports show conditional heteroscedasticity, while the service imports do not. In terms of magnitude, the non-linearity gets stronger after the post-2006 data are included. The estimated η for aggregate imports increases from 112.6 to 168.4 in Model 1. At the same time, the spread between the 95th percentile $\hat{\epsilon}$ and the 5th percentile $\hat{\epsilon}$ increases from 180 percent to 367 percent

($e^{1.541} \approx 4.67$). As noted above, the p-values and the predicted order statistics $\hat{\epsilon}_m$ in Model 2 are not comparable between these two tables. Nevertheless, the point estimate of η is comparable, and it conveys the same message. For aggregate imports in Model 2, the estimated η more than doubles from 0.56 in pre-2006 data to 1.15 when post-2006 data are included.

Model	All 1	All 2	Goods 1	Goods 2	Services 1	Services 2
p^*	2	1	2	1	4	4
k^*	4	2	4	2	$\underline{15}$	$\underline{15}$
$\eta \times 1000$	147.7	0.8844	155.8	0.8547	$\underline{163.3}$	$\underline{0.1248}$
t	3.13	3.128	3.261	3.225	$\underline{1.155}$	$\underline{1.288}$
p-value-bootstrap	0.00413	0.0202	0.004729	0.009303	$\underline{0.1481}$	$\underline{0.1136}$
$\pm\log(\epsilon_{95}/\epsilon_5)$	0.8161	1.029	0.8714	1.074	$\underline{0.2876}$	$\underline{0.292}$
$\pm\log(\epsilon_{90}/\epsilon_{10})$	0.5786	0.5689	0.5992	0.591	$\underline{0.2276}$	$\underline{0.234}$
No.Obs	133	133	133	133	133	133

Model	Services 1	Services 2
p^*	4	4
k^*	16	24
$\eta \times 1000$	232.1	0.231
t	1.56	1.377
p-value-bootstrap	0.128	0.04494
$\pm\log(\epsilon_{95}/\epsilon_5)$	0.4107	0.3726
$\pm\log(\epsilon_{90}/\epsilon_{10})$	0.3276	0.2936
No.Obs	123	123

Table 4.3: U.S. Aggregate Imports, 1970 – 2006, HP Filter

Note: This table reports the estimation of equation (4.1), (4.2), and (4.3) using U.S. aggregate imports data. Data source: NIPA Table 1.1.5. Time frame is 1st quarter of 1970 to the 4th quarter of 2006. The variables of interest are HP-filtered (weight = 1600) import-to-GDP ratios. See the main text for the explanation of the statistics with underlines.

The upper panel in Table 4.3 reports the robustness check using an HP filter for the period between 1970 and 2006. The results are qualitatively similar to those based on the linear filter: both the aggregate and goods imports show non-linearity. The magnitude of non-linearity is higher with an HP filter. For example, the estimated η

for aggregate imports in Model 1 is 147.7 using the HP filter, while it is 112.6 with the linear filter. For the imports of goods, the estimated η for Model 1 increases from 123.2 to 155.8. The results based on Model 2 are similar.

Note that the estimation for the imports of services with the HP filter is problematic, as the optimally chosen k^* hits the upper bound of $k_{max} = 15$ for both models. To ensure an interior solution, we set $k_{max} = 25$ while keeping $p_{max} = 15$ and re-estimate both models. The results are presented in the lower panel of the same table. The point estimate in Model 1 is large. However, the t statistics and the bootstrapped p-value indicate that the large estimate is a result of lower precision. In Model 2 the estimated η is significantly greater than zero with p-value less than 5 percent. However, the magnitude of non-linearity is relatively small: the $\hat{\epsilon}_{95}$ is only 50 percent higher than $\hat{\epsilon}_5$.

4.4.2 Exports

We repeat the same exercise with U.S. aggregate exports data from 1970:1 to 2006:4, and the results are reported in Table 4.4. The results are similar to those based on aggregate imports: we can reject the hypothesis of conditional homoscedasticity in aggregate and goods exports, but we cannot reject the hypothesis for service exports. Overall, the magnitude of non-linearity is smaller in aggregate exports: the estimated η for Model 1 is 34.07, and the 95th/5th percentile spread is 66 percent.

Table 4.5 reports the results based on the data series extended to the end of 2012. Different from the imports data, both the predicted $\hat{\epsilon}^2$ and $|\epsilon|$ are strictly positive for both models. After the inclusion of post-2006 data, the qualitative message does not change. We can still observe conditional heteroscedasticity among aggregate and goods exports, but not among service exports. Also similar to the imports data, post-2006 data increase the magnitude of non-linearity. For example, the estimated η in Model 1 for aggregate exports increases from 34.07 to 46.79, while the 95th/5th

Model	All 1	All 2	Goods 1	Goods 2	Services 1	Services 2
p^*	2	2	3	3	1	1
k^*	4	3	6	5	2	3
$\eta \times 1000$	34.07	0.135	34.6	0.1237	5.92	0.02362
t	2.411	2.638	2.155	2.044	0.255	0.596
p-value-bootstrap	0.03591	0.03831	0.05146	0.06563	0.4457	0.3578
$\pm\log(\epsilon_{95}/\epsilon_5)$	0.5102	0.5494	0.5508	0.6106	0.0639	0.1721
$\pm\log(\epsilon_{90}/\epsilon_{10})$	0.4124	0.4509	0.3913	0.4232	0.05198	0.1414
No.Obs	133	133	133	133	133	133

Table 4.4: U.S. Aggregate Exports, 1970 – 2006, Linear Filter

Note: This table reports the estimation of equation (4.1), (4.2), and (4.3) using U.S. aggregate exports data. Data source: NIPA Table 1.1.5. Time frame is 1st quarter of 1970 to the 4th quarter of 2006. The variables of interest are linearly-filtered export-to-GDP ratios.

spread in Model 1 increases from 66 percent to 101 percent.

Model	All 1	All 2	Goods 1	Goods 2	Services 1	Services 2
p^*	2	2	2	2	1	1
k^*	4	6	4	6	3	3
$\eta \times 1000$	46.79	0.2403	41.16	0.2171	20.14	0.03432
t	3.067	2.113	2.349	1.784	1.009	1.056
p-value-bootstrap	0.004467	0.02623	0.03318	0.04727	0.2687	0.2278
$\pm\log(\epsilon_{95}/\epsilon_5)$	0.7	0.9017	0.588	0.7644	0.233	0.2745
$\pm\log(\epsilon_{90}/\epsilon_{10})$	0.616	0.7012	0.5018	0.6307	0.1845	0.2127
No.Obs	157	157	157	157	157	157

Table 4.5: U.S. Aggregate Exports, 1970 – 2012, Linear Filter

Note: This table reports the estimation of equation (4.1), (4.2), and (4.3) using U.S. aggregate exports data. Data source: NIPA Table 1.1.5. Time frame is 1st quarter of 1970 to the 4th quarter of 2012. The variables of interest are linearly-filtered export-to-GDP ratios.

The upper panel of Table 4.6 reports the results from 1970 to 2006 based on the HP filter. In this case the aggregate exports-to-GDP ratio exhibits conditional heteroscedasticity, while the estimated η for the exports in services is not significantly different from zero. The estimation for the exports in goods is problematic, because the optimally chosen $k^* = k_{max} = 15$. Again, we re-estimate the model with $k_{max} = 25$

and report the results in the lower panel. The point estimates for both models are large, however, they are not significantly different from zero at the 5 percent level.

Model	All 1	All 2	Goods 1	Goods 2	Services 1	Services 2
p^*	2	4	4	4	1	1
k^*	5	7	15	15	3	4
$\eta \times 1000$	85.37	0.2135	46.61	0.1685	79.14	0.1233
t	2.937	2.289	0.8627	1.063	1.217	1.289
p-value-bootstrap	0.009956	0.07102	0.2019	0.146	0.2396	0.1711
$\pm \log(\epsilon_{95}/\epsilon_5)$	0.5837	0.4189	0.2021	0.2268	0.3269	0.38
$\pm \log(\epsilon_{90}/\epsilon_{10})$	0.4763	0.3408	0.1687	0.1904	0.2366	0.2764
No.Obs	133	133	133	133	133	133

Model	Goods 1	Goods 2
p^*	3	3
k^*	20	20
$\eta \times 1000$	101.5	0.3427
t	1.286	1.46
p-value-bootstrap	0.07526	0.1035
$\pm \log(\epsilon_{95}/\epsilon_5)$	0.3529	0.4073
$\pm \log(\epsilon_{90}/\epsilon_{10})$	0.2521	0.2813
No.Obs	123	123

Table 4.6: U.S. Aggregate Exports, 1970 – 2006, HP Filter

Note: This table reports the estimation of equation (4.1), (4.2), and (4.3) using U.S. aggregate exports data. Data source: NIPA Table 1.1.5. Time frame is 1st quarter of 1970 to the 4th quarter of 2006. The data is filtered with an HP filter with weight 1600. See the main text for the explanation of the statistics with underlines.

4.4.3 Real Exchange Rates

In the above two sections we document that both the aggregate U.S. imports and exports show strong conditional heteroscedasticity. In this section we document that the price of U.S. dollars against a trade-weighted currency index exhibit a similar pattern. The main results are reported in Tables 4.7 and 4.8.

Unlike the imports and exports data, the real dollar index is stationary for the

(a) Unfiltered, 2006

Model	1	2
p^*	2	2
k^*	10	9
$\eta \times 1000$	57.85	230.3
t	3.264	2.821
p-value-bootstrap	0.00807	0.0259
$\pm\log(\epsilon_{95}/\epsilon_5)$	0.5986	0.439
$\pm\log(\epsilon_{90}/\epsilon_{10})$	0.5223	0.3901
No.Obs	121	121

(b) Linear Filter, 2006

Model	1	2
p^*	2	2
k^*	10	9
$\eta \times 1000$	57.02	224.2
t	3.225	2.753
p-value-bootstrap	0.008705	0.0295
$\pm\log(\epsilon_{95}/\epsilon_5)$	0.576	0.4182
$\pm\log(\epsilon_{90}/\epsilon_{10})$	0.5213	0.3792
No.Obs	121	121

Table 4.7: U.S. Real Exchange Rate, 1973 - 2006

Note: This table reports the estimation of equation (4.1), (4.2), and (4.3) using real trade-weighted U.S. dollar index. Data source: Board of Governors, Federal Reserve Bank.

period considered. Therefore, the benchmark results presented in the left panels of both tables are based on the unfiltered data. Nevertheless, we report the linearly-filtered data in the right panels of the same tables.

Table 4.7 reports the results based on data from 1973 to 2006. The real exchange rate under both models shows significant conditional heteroscedasticity. With both filtered and unfiltered data, the estimated η are significantly different from zero. The magnitude of non-linearity is also substantial. For example, with unfiltered data, the 95th/5th percentile spread of $\hat{\epsilon}$ is 82 percent in Model 1 and 55 percent in Model 2. The differences between the unfiltered and linearly-filtered data are small. For example, the point estimates of η in Model 1 only differ by 1.4 percent (57.85 v.s. 57.02); they differ by 2.7 percent in Model 2 (230.3 v.s. 224.2).

Table 4.8 reports the results based on the data from 1973 to 2012. Again, the real exchange rate shows strong conditional heteroscedasticity in both models. Different from imports and exports, the inclusion of post-2006 data does not change the results by a large margin. For example, the point estimate of η changes from 57.85 to 54.79 in Model 1 with unfiltered data. Similar results can be found with Model 2. Once again, the differences between the unfiltered and linearly-filtered data are small.

(a) Unfiltered, 2012

Model	1	2
p^*	2	2
k^*	12	12
$\eta \times 1000$	54.79	231.8
t	3.296	3.01
p-value-bootstrap	0.00667	0.0295
$\pm \log(\epsilon_{95}/\epsilon_5)$	0.5452	0.4239
$\pm \log(\epsilon_{90}/\epsilon_{10})$	0.4675	0.3676
No.Obs	145	145

(b) Linear Filter, 2012

Model	1	2
p^*	2	2
k^*	10	9
$\eta \times 1000$	52.43	201.1
t	3.148	2.637
p-value-bootstrap	0.00990	0.05301
$\pm \log(\epsilon_{95}/\epsilon_5)$	0.5088	0.37
$\pm \log(\epsilon_{90}/\epsilon_{10})$	0.4684	0.3324
No.Obs	145	145

Table 4.8: U.S. Real Exchange Rate, 1973 - 2012

Note: This table reports the estimation of equation (4.1), (4.2), and (4.3) using real trade-weighted U.S. dollar index. Data source: Board of Governors, Federal Reserve Bank.

4.5 Conclusion

We provide empirical evidence showing that U.S. aggregate imports, exports, and real exchange rate data exhibit strong conditional heteroscedasticity. In an autoregression framework, this means that the responses of variables to shocks are history-dependent. The response will be stronger when the past values of the variables are higher. We document this empirical pattern for imports and exports from 1970 to 2006. For most of our specifications, conditional heteroscedasticity is more pronounced for trade in goods than for trade in services.

This novel empirical pattern is interesting because it runs at odds with a standard IRBC model, where the impact responses of imports, exports, and real exchange rate are not history-dependent. We speculate that a trade model with inventories shall be able to explain these new empirical patterns, and the intuition is as follows.

Suppose we have a two-country model, where the firms in the home country need to import from the foreign country, and then sell to the domestic consumers. Also assume that the importers need to hold inventories of the foreign-produced goods. In a dynamic setting, the distribution of inventories across firms will be endogenous.

After an expansion — a sequence of positive shocks — firms will, on average, deplete their inventories of foreign goods due to high domestic demand. A subsequent positive shock at the end of an expansion is then likely to induce a large response in imports, because many firms need to replenish their inventory holdings. On the other hand, after a prolonged recession — a sequence of negative shocks — firms will, on average, hold high levels of inventories due to weak domestic demand. Another positive shock at the end of a recession is then likely to induce a much smaller response in imports, because many firms can rely on the stockpiled inventories to fulfill the surge in demand. In this way, the impulse response of aggregate imports will be history-dependent, just as we documented in this paper.

In a symmetric world, the importers in the foreign country also need to hold inventories. Subsequently, the imports in the foreign country will also exhibit conditional heteroscedasticity. This implies that the exports in the home country will follow a similar pattern and show conditional heteroscedasticity as well. The real exchange rate will also likely show conditional heteroscedasticity, because the responses of trade volumes are much higher at the end of an expansion relative to the end of a recession.

APPENDIX

APPENDIX A

A.1 Details of the Model

A.1.1 The Firm's Problem

Denote the total expenditure in country i as H_i, the ideal price level as P_i. If a firm in country j wants to sell to the market i, denote the price of the good as $p_{ij}(x)$ and the marginal cost (iceberg cost included) of selling to market i as $M_{ij}(x)$. The firm solves the following problem:

$$\begin{aligned} \max_{q_{ij}(x)} \quad & p_{ij}(x)q_{ij}(x) - M_{ij}(x)q_{ij}(x), \\ \text{s.t.} \quad & p_{ij}(x) = H_i^{\frac{1}{\epsilon}} P_i^{\frac{\epsilon-1}{\epsilon}} q_{ij}(x)^{-\frac{1}{\epsilon}}, \end{aligned}$$

where the constraint of the maximization problem is the inverse of the derived demand function from solving the consumer's problem in market i.

The solution of the above maximization problem is

$$q_{ij}(x) = H_i P_i^{\epsilon-1} \left(\frac{\epsilon}{\epsilon-1} M_{ij}(x) \right)^{-\epsilon}, \tag{A.1}$$

$$p_{ij}(x) = \frac{\epsilon}{\epsilon-1} M_{ij}(x). \tag{A.2}$$

Equation (A.2) is the result of plugging equation (A.1) into the inverse derived demand function.

The marginal cost of supplying to market i depends on the productivity of the firm, as well as the method through which the firm chooses to serve market i. If market i is served by a domestic firm or by an exporter in country j, then:

$$M_{ij}(x) = \frac{\tau_{ij} w_j}{A_j(x)}.$$

In the special case of $i = j$, market i is served by the domestic firm in country i:

$$M_{ii}(x) = \frac{w_i}{A_i(x)}.$$

If market i is served by an MNE founded in country j, then

$$M_{ij}(x) = \frac{w_i}{A_j(x)}.$$

The sales to market i, $\sigma_{ij}(x)$ is therefore

$$\sigma_{ij}(x) = p_{ij}(x) q_{ij}(x) = H_i P_i^{\epsilon-1} \left(\frac{\epsilon}{\epsilon-1} M_{ij}(x) \right)^{1-\epsilon}.$$

To supply $q_{ij}(x)$ to market i, the labor used in production is

$$L_{ij}(x) = H_i P_i^{\epsilon-1} \left(\frac{\epsilon}{\epsilon-1} M_{ij}(x) \right)^{\epsilon} \frac{\tau_{ij}}{A_j(x)},$$

with the understanding that when $i = j$, $\tau_{ij} = 1$.

The profit earned in market i before the fixed cost is

$$p_{ij}(x) - M_{ij}(x)]q_{ij}(x) = \frac{H_i}{\epsilon} P_i^{\epsilon-1} \left(\frac{\epsilon}{\epsilon - 1} M_{ij}(x) \right)^{1-\epsilon}$$

To ensure that firms sort into non-exporters, exporters, and multinational firms by productivity, I impose the following assumption similar to the one used in *Helpman et al.* (2004):

$$\frac{g_{ji}}{f_{ji}} \leq \left(\frac{\tau_{ji} w_i}{w_j} \right)^{\epsilon-1}$$

This equation implies that only the most productive firms will engage in FDI, while the other productive firms choose export over FDI.

A similar restriction needs to be imposed to ensure the separation of the domestic firms: we need to make sure that in equilibrium, not all the firms choose to sell to the foreign market. In a Melitz model, this condition can be written down explicitly. Unfortunately, it is not possible to do so for this paper. The reason is that x_i^* does not admit a closed-form solution. Nevertheless, characterization of the restriction is still possible. Generally, we need the market size of the home country to be above a certain level relative to the foreign country, or the variable trade cost to be above a certain level, so the firms in the home country will not find exporting to the foreign country too easy. In all the results presented in this paper, the separation of firms into domestic and exporting/multinational firms is checked and ensured.

A.1.2 The Equilibrium Conditions

The first three equilibrium conditions on cutoff human capital levels are self-evident. Here I explain the other two equilibrium conditions in detail. In this section, I derive the equilibrium conditions under truncation.

Income-Expenditure Identity The third equilibrium condition, equation (2.8), requires that the total expenditure and total income in country i must be the same. Total expenditure is denoted as H_i. Total income consists of two parts: the total labor income and the total profits. The CEO compensation function, $k(\pi)$, does not enter the accounting equation. The difference between the profit and the CEO compensation at each firm is distributed to all the individuals in the same country, and therefore $k(\pi)$ does not matter for total income.

The total labor income is easy to compute. It is the wage rate $w(i)$ times the total labor supply:

$$w_i \cdot \left(n_i \int_0^{x_i^*} x f_i(x) dx \right) = w_i n_i \frac{\lambda}{1 - e^{s_i \lambda}} \int_0^{x_i^*} x e^{-\lambda x} dx, \tag{A.3}$$

$$= \frac{w_i n_i}{(1 - e^{-\lambda s_i} \lambda)} \left[e^{-\lambda x_i^*} (-\lambda x_i^* - 1) + 1 \right], \tag{A.4}$$

$$= w_i n_i \frac{F(x_i^*)}{\lambda} - \frac{n_i x_i^* e^{-\lambda x_i^*}}{1 - e^{-\lambda s_i}}, \tag{A.5}$$

$$= w_i \cdot \underbrace{\left\{ \frac{n_i}{\lambda} [F(x_i^*) - x_i^* f(x_i^*)] \right\}}_{\text{Labor Supply}}, \tag{A.6}$$

where $f(.)$ is the PDF of the truncated exponential distribution. The part in the curly brackets is the total labor supply in country i.

The total profit in country i is composed of three parts: the profit earned in the home country i, the profit earned in the other country j through export, and the profit earned in country j through FDI. This three-part separation is not the same as separating the profits into firms in the three corresponding groups. The difference is that, the profits earned in the home country i includes the profits from all the firms, as the exporters and MNEs also sell to the home market.

The total profit earned in the home market i is

$$n_i \int_{x_i^*}^{s} \frac{H_i}{\epsilon} P_i^{\epsilon-1} \left(\frac{\epsilon}{\epsilon-1} w_i\right)^{1-\epsilon} (b_i e^x)^{\epsilon-1} f_i(x) dx - n_i f_{ii} w_i [1 - F(x_i^*)].$$

The total profit earned in the foreign market though exporting is

$$n_i \int_{x_{ji}^e}^{x_{ji}^f} \frac{H_j}{\epsilon} P_j^{\epsilon-1} \left(\frac{\epsilon}{\epsilon-1} \tau_{ji} w_i\right)^{1-\epsilon} (b_i e^x)^{\epsilon-1} f_i(x) dx - n_i f_{ji} w_i [F(x_{ji}^f) - F(x_{ji}^e)],$$

and the total profit earned in the foreign market through FDI is

$$n_i \int_{x_{ji}^f}^{s} \frac{H_j}{\epsilon} P_j^{\epsilon-1} \left(\frac{\epsilon}{\epsilon-1} w_j\right)^{1-\epsilon} (b_i e^x)^{\epsilon-1} f_i(x) dx - n_i g_{ji} w_i [1 - F(x_{ji}^f)].$$

The total profit in country i is the summation over these three parts. The income-expenditure identity here does not imply trade balance, as it usually does in a Melitz model. What it does imply is trade and financial balance. Trade in equilibrium is almost surely unbalanced, and the gap will be offset by the differences in capital flow: the differences between the profits the domestic MNEs collected from abroad and the foreign MNEs collected from the home market.

Ideal Price Level Equation (2.9) is the definition of the ideal price level in country i. What needs further explanation is the set of goods available in country i: Θ_i. This set is the union of three mutually exclusive subsets: (1) the goods provided by all the firms created in country i, (2) the goods provided by all the exporting firms in country j, and (3) the goods provided by all the MNEs in country j. The price for every single variety in each of the subsets is a constant mark-up over the marginal cost in that subset. The marginal cost for goods in different subsets can be found in Appendix A.1.1. The ideal price level is a CES integration of all the individual prices over the set Θ_i.

After decomposing the set Θ_i into the three subsets mentioned above, the ideal price level can be expressed based on the firm productivity distribution directly:

$$P_i^{1-\epsilon} = \left\{ \sum_{j=1}^{2} \left[n_j \left(\frac{\epsilon}{\epsilon - 1} \tau_{ij} w_j \right)^{1-\epsilon} \int_{x_{ij}^e}^{x_{ij}^f} b_i e^x f(x) dx + n_j \left(\frac{\epsilon}{\epsilon - 1} w_i \right)^{1-\epsilon} \int_{x_{ij}^f}^{s} b_i e^x f(x) dx \right] \right\}.$$

Note that when $i = j$, $x_{ij}^e = x_i^*$. The first part in the square bracket includes all the goods provided by domestic firms, domestic exporters, and foreign exporters. The second part in the square bracket includes all the goods provided by the domestic and foreign MNEs.

A.1.3 Firm Size Distributions

In this appendix, I derive the CDF of firm productivity, sales, profit, and employment distributions for different groups of firms.

A.1.3.1 Productivity Distribution

The human capital, x, in country i is distributed exponentially with the following CDF:

$$F(x) = 1 - e^{-\lambda x},$$

and the firm founded by the individual with human capital x has the following productivity:

$$A_i(x) = b_i e^x.$$

The CDF of the firm productivity distribution, denoted as $F_A(y)$, can be derived as follows:

$$F_A(y) = \Pr(A_i(x) \leq y) = \Pr(b_i e^x \leq y) = \Pr(e^x \leq \frac{y}{b_i}),$$
$$= \Pr(x \leq \log(y/b_i)) = F(\log(y/b_i)),$$
$$= 1 - e^{-\lambda \log(y/b_i)},$$
$$= 1 - b_i^{\lambda} y^{-\lambda},$$

which is the CDF of a Type-I Pareto distribution with location parameter b_i and shape parameter λ. This CDF is shared by all the firms in country i whether they are non-exporting firms, exporting firms, or multinational firms.

Truncation If the exponential distribution is truncated from above at s, then the CDF of the human capital distribution will be

$$F(x) = \frac{1 - e^{-\lambda x}}{1 - e^{-\lambda s}}, x \in [0, s].$$

Given the same functional form of firm productivity, the CDF of the productivity distribution can be derived using similar methods outlined above. The distribution can be verified to be a truncated Pareto distribution,

$$F_A(y) = \frac{1 - b_i^{\lambda} y^{-\lambda}}{1 - b_i^{\lambda} u_i^{-\lambda}}, y \in [b_i, u_i],$$

where u_i is the country-specific upper bound of firm productivity:

$$u_i = b_i e^s.$$

In the rest of the this appendix, I use the original distribution without truncation.

A.1.3.2 Sales Distribution

The sales from country j to country i is derived in Appendix A.1.1 and repeated here:

$$p_{ij}(x)q_{ij}(x) = H_i P_i^{\epsilon-1}\left(\frac{\epsilon}{\epsilon-1}M_{ij}(x)\right)^{1-\epsilon}, \tag{A.7}$$

where $M_{ij}(x)$ is the marginal cost of production conditional on the mode of access (export or multinational production). Based on the market-specific sales, I derive the firm sales. I denote sales for a firm with CEO human capital x in country i as $\sigma_i(x)$ and rewrite it as a linear function of $A_i(x)^{\epsilon-1}$:

$$\sigma_i(x) = \Sigma_i(x)A_i(x)^{\epsilon-1}.$$

$\Sigma_i(x)$ summarizes the market size accessible to the firm. It is a step function depending on x:

$$\Sigma_i(x) = \begin{cases} H_i\left(\frac{P_i}{w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} & , x \in [x_i^*, x_{ji}^e), \\ H_i\left(\frac{P_i}{w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} + H_j\left(\frac{P_j}{\tau_{ji}w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} & , x \in [x_{ji}^e, x_{ji}^f), \\ H_i\left(\frac{P_i}{w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} + H_j\left(\frac{P_j}{w_j}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} & , x \in [x_{ji}^f, \infty). \end{cases}$$

The first line is the market accessible to the non-exporters, the second line the exporters, and the last line the multinational producers. The general formula for the

CDF of the sales distribution is

$$
\begin{aligned}
F_\sigma(y) &= \Pr(\sigma < y), \\
&= \Pr(\Sigma_i(x) A_i(x)^{\epsilon-1} < y) = \Pr\left(A_i(x) < \left(\frac{y}{\Sigma_i(x)}\right)^{\frac{1}{\epsilon-1}} \right), \\
&= F_A\left(\left(\frac{y}{\Sigma_i(x)}\right)^{\frac{1}{\epsilon-1}} \right) = 1 - b_i^\lambda \left(\frac{y}{\Sigma_i(x)}\right)^{\frac{-\lambda}{\epsilon-1}}, \\
&= 1 - \left(\frac{\Sigma_i(x)}{b_i^{1-\epsilon}}\right)^\theta y^{-\theta},
\end{aligned}
$$

where

$$
\theta = \frac{\lambda}{\epsilon - 1}.
$$

The above equation defines Type-I Pareto distribution with shape parameter $\frac{\lambda}{\epsilon-1}$ and location parameter $\Sigma_i(x) b_i^{\epsilon-1}$. The location parameter differs by $\Sigma_i(x)$. The non-exporting firms have the smallest $\Sigma_i(x)$ and therefore the lowest location parameter. The exporting firms have higher $\Sigma_i(x)$ and the multinational firms have the highest $\Sigma_i(x)$. Note that within the same group (non-exporters, exporters, and multinationals), $\Sigma_i(x)$ is the same for all the firms.

A.1.3.3 Profit Distribution

The profit earned in each market is provided in Appendix A.1.1. Based on the market-specific profit, the firm profit can be written as an affine function of $A_i(x)^{\epsilon-1}$:

$$
\pi_i(x) = \Pi_i(x) A_i(x)^{\epsilon-1} - C_i(x).
$$

Similar to the sales distribution, $\Pi_i(x)$ takes three values depending on x:

$$\Pi_i(x) = \begin{cases} \frac{H_i}{\epsilon}\left(\frac{P_i}{w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} & , x \in [x_i^*, x_{ji}^e), \\ \frac{H_i}{\epsilon}\left(\frac{P_i}{w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} + \frac{H_j}{\epsilon}\left(\frac{P_j}{\tau_{ji}w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} & , x \in [x_{ji}^e, x_{ji}^f), \\ \frac{H_i}{\epsilon}\left(\frac{P_i}{w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} + \frac{H_j}{\epsilon}\left(\frac{P_j}{w_j}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} & , x \in [x_{ji}^f, \infty). \end{cases}$$

The first line is the market size accessible to a domestic firm. The second line is the market size for exporting firms, and the third line is the market size for multinational firms. Similarly, the fixed cost term $C_i(x)$ depends on the type of the firm

$$C_i(x) = \begin{cases} w_i f_{ii} & , x \in [x_i^*, x_{ji}^s), \\ w_i(f_{ii} + f_{ji}) & , x \in [x_{ji}^e, x_{ji}^f), \\ w_i(f_{ii} + g_{ji}) & , x \in [x_{ji}^f, \infty). \end{cases}$$

The distribution function of π takes the following general formula

$$\begin{aligned} F_\pi(y) = \Pr(\pi \leq y) &= \Pr(\Pi_i(x) \cdot A_i(x)^{\epsilon-1} - C_i(x) \leq y), \\ &= \Pr\left(A_i(x) \leq \left(\frac{y + C_i(x)}{T_i(x)})\right)^{\frac{1}{\epsilon-1}}\right), \\ &= 1 - b_i^\lambda \left(\frac{y + C_i(x)}{T_i(x)})\right)^{\frac{-\lambda}{\epsilon-1}} = 1 - \left(\frac{y + C_i(x)}{T_i(x) b_i^{\epsilon-1}}\right)^{-\frac{\lambda}{\epsilon-1}}, \\ &= 1 - \left(1 + \frac{y + \mu_i(x)}{\chi_i(x)}\right)^{-\theta}, \end{aligned}$$

where

$$\begin{aligned} \mu_i(x) &= \chi_i(x) - C_i(x), \\ \chi_i(x) &= T_i(x) \cdot b_i^{\epsilon-1}, \\ \theta &= \frac{\lambda}{\epsilon - 1}. \end{aligned}$$

This equation is the CDF of a Type-II Pareto distribution as defined in *Arnold* (1985).

The shape index of the firm profit distribution is $\theta = \frac{\lambda}{\epsilon-1}$. The two location parameters $\mu_i(x)$ and $\chi_i(x)$ depend on the market that the firm can access to.

A.1.3.4 Employment Distribution

Employment distribution is similar to the profit distribution. Market-specific employment is provided in Appendix A.1.1 and here I aggregate it up to firm-level employment. For each firm the employment, $L_i(x)$, can be written as an affine function of $A_i(x)^{\epsilon-1}$:

$$L_i(x) = \Lambda_i(x) A_i(x)^{\epsilon-1} + T_i(x).$$

$\Lambda_i(x)$, again, summarizes the market size accessible to a firm x and is a step function that takes three values:

$$\Lambda_i(x) = \begin{cases} \frac{H_i}{P_i^{1-\epsilon}} \left(\frac{1}{w_i} \frac{\epsilon-1}{\epsilon} \right)^{\epsilon} & , x \in [x_i^*, x_{ji}^e), \\ \frac{H_i}{P_i^{1-\epsilon}} \left(\frac{1}{w_i} \frac{\epsilon-1}{\epsilon} \right)^{\epsilon} + \frac{H_j}{P_j^{1-\epsilon}} \left(\frac{1}{w_i} \frac{\epsilon-1}{\epsilon} \right)^{\epsilon} \tau_{ji}^{1-\epsilon} & , x \in [x_{ji}^e, x_{ji}^f), \\ \frac{H_i}{P_i^{1-\epsilon}} \left(\frac{1}{w_i} \frac{\epsilon-1}{\epsilon} \right)^{\epsilon} + \frac{H_j}{P_j^{1-\epsilon}} \left(\frac{1}{w_j} \frac{\epsilon-1}{\epsilon} \right)^{\epsilon} & , x \in [x_{ji}^f, \infty). \end{cases}$$

$T_i(x)$ is the labor used as fixed cost of operation, export, and multinational production:

$$T_i(x) = \begin{cases} f_{ii} & , x \in [x_i^*, x_{ji}^s), \\ f_{ii} + f_{ji} & , x \in [x_{ji}^e, x_{ji}^f), \\ f_{ii} + g_{ji} & , x \in [x_{ji}^f, \infty). \end{cases}$$

Because both the employment and the profit are affine transformations of $A_i(x)^{\epsilon-1}$, the steps to derive the general formula of CDF are exactly the same. In the end, employment distributions are also Type-II Pareto distributions with shape parameter θ. The two location parameters depend on the market size accessible to the firm as well.

A.1.4 Income Distribution

The equilibrium income distribution in the model follows a two-class structure: the worker's income distribution follows an exponential distribution, and the CEO's income follows various Pareto-Type distributions. In this appendix, I present the details of the income distributions of the model.

Workers Workers in country i receive w_i for each unit of efficiency labor supplied to the market. The income for a worker with human capital x is $w_i x$, which follows an exponential distribution, same as x. The shape parameter of the income distribution is $\frac{\lambda}{w_i}$. The CDF of the distribution is

$$\begin{aligned} V(y) &= \Pr(w_i x \leq y) = \Pr(x \leq \frac{y}{w_i}), \\ &= 1 - e^{-\frac{\lambda}{w_i} y}. \end{aligned}$$

CEOs If $k(\pi)$ is monotonic and regularly varying with tail index β, then the CEO income follows a Pareto-Type distribution with shape parameter θ/β. Given a compensation function $k(\pi)$, the CDF of the CEO income is

$$U(y) = \Pr(k(\pi) \leq y) = \Pr(\pi \leq k^{-1}(y)) = F_\pi(k^{-1}(y)),$$

where $k^{-1}(y)$ is the inverse of $k(\pi)$ and $F_\pi(\cdot)$ is the CDF of firm profit distribution derived in Appendix A.1.3. The inverse function exists because $k(\pi)$ is monotonic. Because $k(\pi)$ is a regularly varying function with tail index β, the inverse function $k^{-1}(\cdot)$ is also a regularly varying function with tail index $1/\beta$ (Proposition 0.8.5, *Resnick* (1987)).

The survival function of π is a regularly varying function, with tail index $-\theta$ as well. To see this:

$$\lim_{\pi\to\infty}\frac{1-F_\pi(\eta\pi)}{1-F_\pi(\pi)} = \frac{\left(1+\frac{\eta\pi+\mu}{\chi}\right)^{-\theta}}{\left(1+\frac{\pi+\mu}{\chi}\right)^{-\theta}} = \eta^{-\theta}.$$

The composition of two regularly varying functions is a regularly varying function, and the tail index of the composition function is the product of the two indices (Proposition 0.8.4, *Resnick* (1987)). Therefore $1-U(y)$, as the composition of $k^{-1}(y)$ and $1-F_\pi(\pi)$, is a regularly varying function with tail index $-\frac{\theta}{\beta}$. This defines $y=k(\pi)$ as a Pareto-Type distribution with shape parameter $\frac{\theta}{\beta}$ (Definition 7.25, *Gulisashvili* (2012)). Moreover, the CDF of $k(\pi)$ can be re-written as:

$$U(y) = 1 - y^{-\theta/\beta}R(y),$$

where $R(y)$ is a slowly varying function:

$$\lim_{y\to\infty}\frac{R(\eta y)}{R(y)} = 1.$$

Example The CEO compensation function for corporations defined in Section 2.5 is

$$k(\pi) = \alpha^{1-\beta}\pi^{\beta} = \alpha^{1-\beta}\left(\Pi\cdot A^{\epsilon-1}-C\right)^{\beta}.$$

The CDF of $k(\pi)$ is

$$
\begin{aligned}
U(y) = \Pr(k \leq y) &= \Pr(\alpha^{1-\beta}\left(\Pi \cdot A^{\epsilon-1} - C\right)^{\beta} \leq y), \\
&= \Pr\left(A^{\epsilon-1} \leq \frac{y^{\frac{1}{\beta}}\alpha^{\frac{\beta-1}{\beta}} + C}{\Pi}\right), \\
&= 1 - b^{\lambda}\left(\frac{y^{\frac{1}{\beta}}\alpha^{\frac{\beta-1}{\beta}} + C}{\Pi}\right)^{-\frac{\lambda}{\epsilon-1}}.
\end{aligned}
$$

Using the general result proved above, it is trivial to show that $k(\pi)$ follows a Pareto-Type distribution. Here I follow a different route and prove directly that the survival function $1 - U(y)$ is a regularly varying function. To see this:

$$
\begin{aligned}
\lim_{y\to\infty}\frac{1-U(\eta y)}{1-U(y)} &= \lim_{y\to\infty}\left(\frac{\eta^{\frac{1}{\beta}}y^{\frac{1}{\beta}}\alpha^{\frac{\beta-1}{\beta}} + C}{y^{\frac{1}{\beta}}\alpha^{\frac{\beta-1}{\beta}} + C}\right)^{-\frac{\lambda}{\epsilon-1}}, \\
&= \lim_{y\to\infty}\left(\frac{\eta^{\frac{1}{\beta}} + \frac{C}{y^{\frac{1}{\beta}}\alpha^{\frac{\beta-1}{\beta}}}}{1 + \frac{C}{y^{\frac{1}{\beta}}\alpha^{\frac{\beta-1}{\beta}}}}\right)^{-\frac{\lambda}{\epsilon-1}}.
\end{aligned}
$$

As $y \to \infty$, $y^{\frac{1}{\beta}} \to \infty$, therefore

$$
\lim_{y\to\infty}\frac{1-U(\eta y)}{1-U(y)} = \eta^{-\frac{\lambda}{\beta(\epsilon-1)}},
$$

which defines $1 - U(y)$ as a regularly varying function with index $-\frac{\lambda}{\beta(\epsilon-1)}$. This further implies that the income distribution function of CEOs in corporations can be expressed as

$$
U(y) = 1 - y^{-\frac{\lambda}{\beta(\epsilon-1)}}R(y).
$$

The income distribution of the CEOs at sole proprietorship firms is the same as the profit distribution and therefore is Type-II Pareto.

See *Feller* (1966), *Resnick* (1987), and *Gulisashvili* (2012) for more details on regularly varying functions and Pareto-Type distributions.

A.1.5 Profit-to-Wage Ratios

Profit-to-wage ratios in this model only depends on the cutoff human capitals in general equilibrium. This property can be exploited to gain some insight into the basic mechanism of the model without quantification.

Domestic Profit The profit-to-wage ratio in the domestic market is the profit earned from the domestic market divided by domestic wage. This part of profit is earned by the domestic firms, the exporters, and the MNEs created in the home country.

The profit-to-wage ratio is

$$\frac{\pi_{ii}(x)}{w_i} = \frac{H_i}{w_i \epsilon} \left(\frac{P_i}{w_i} \frac{\epsilon - 1}{\epsilon} \right)^{\epsilon - 1} A_i(x)^{\epsilon - 1} - f_{ii}.$$

From the cutoff condition of the marginal firm, we know:

$$\frac{H_i}{w_i \epsilon} \left(\frac{P_i}{w_i} \frac{\epsilon - 1}{\epsilon} \right)^{\epsilon - 1} b_i^{\epsilon - 1} e^{(\epsilon - 1) x_i^*} - f_{ii} = x_i^*,$$

and therefore

$$\frac{H_i}{w_i \epsilon} \left(\frac{P_i}{w_i} \frac{\epsilon - 1}{\epsilon} \right)^{\epsilon - 1} = \frac{x_i^* + f_{ii}}{b_i e^{(\epsilon - 1) x_i^*}}. \tag{A.8}$$

Plug this into the first equation, we have

$$\frac{\pi_{ii}(x)}{w_i} = (x_i^* + f_{ii}) e^{(\epsilon - 1)(x - x_i^*)} - f_{ii}.$$

The partial derivative of this ratio with respect to x is positive, so in general, the profit-to-wage ratio is higher when the firm is more productive and larger. All the general equilibrium movements affect this ratio through the only endogenous variable in this equation: the cutoff value x_i^*. The cutoff human capital is a measure of the competitiveness of the home market in general equilibrium: it will be higher when the market is more competitive due to highly productive foreign firms entering. The partial derivative of this ratio with respect to x_i^* is

$$\frac{\partial}{\partial x_i^*}\left(\frac{\pi_{ii}(x)}{w_i}\right) = e^{(\epsilon-1)(x-x_i^*)}[1-(\epsilon-1)(x_i^*+f_{ii})]. \tag{A.9}$$

The sign of this derivative is the same as $[1-(\epsilon-1)(x_i^*+f_{ii})]$. I claim that this sign is always negative under the assumption that the least productive individual in country i must not find creating a new firm profitable. This restriction is imposed to guarantee the existence and uniqueness of the occupational choice cutoff in Section 2.3. This assumption means:

$$\begin{aligned} \frac{H_i}{\epsilon}P_i^{\epsilon-1}w_i^{1-\epsilon}\left(\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1} A_i(0)^{\epsilon-1} - f_{ii}w_i < 0, \\ f_{ii} > \frac{H_i}{\epsilon w_i}^{\epsilon-1}\left(\frac{\epsilon-1}{\epsilon}\frac{P_i}{w_i}\right)^{\epsilon-1} A_i(0)^{\epsilon-1}. \end{aligned}$$

Plug equation (A.8) into the above inequality, we have

$$\begin{aligned} f_{ii} &> \frac{x_i^*+f_{ii}}{A_i(x_i^*)^{\epsilon-1}}A_i(0)^{\epsilon-1} \\ f_{ii} &> \frac{x_i^*}{e^{(\epsilon-1)x_i^*}-1}. \end{aligned}$$

Now I need to prove

$$x_i^* + f_{ii} > \frac{1}{\epsilon-1}. \tag{A.10}$$

To do this, I define

$$m(x_i^*) = x_i^* + \frac{x_i^*}{e^{(\epsilon-1)x_i^*} - 1} - \frac{1}{\epsilon - 1}.$$

It is easy to show that $m(x_i^*)$ is monotonically increasing,

$$\frac{\partial m(x_i^*)}{\partial x_i^*} = 1 + \frac{e^{(\epsilon-1)x_i^*}(1 + (\epsilon-1)x_i^*) - 1}{(e^{(\epsilon-1)x_i^*} - 1)^2} > 0,$$

because

$$((\epsilon - 1)x_i^* > 0) \wedge (e^{(\epsilon-1)x_i^*} > 1).$$

Therefore, the minimum of $m(x^*)$ is obtained at $x_i^* = 0$, which is precisely 0. To see this, we need to apply L'Hôpital's rule to the second term at $x_i^* = 0$:

$$\begin{aligned} \lim_{x_i^* \to 0} m(x^*) &= x_i^* + \frac{1}{e^{(\epsilon-1)x_i^*}(\epsilon - 1)} - \frac{1}{\epsilon - 1}, \\ &= \frac{1}{\epsilon - 1} - \frac{1}{\epsilon - 1} = 0. \end{aligned}$$

This implies that for all possible values of $x_i^* \in [0, \infty)$, equation (A.10) is true and therefore the profit-to-wage ratio decreases with x_i^*.

Exporting Profits The profits earned from exporting to the foreign country, divided by local wage, is

$$\frac{\pi_{ji}^e(x)}{w_i} = \frac{H_j}{w_i \epsilon} \left(\frac{P_j}{\tau_{ji} w_i} \frac{\epsilon - 1}{\epsilon} \right)^{\epsilon - 1} A_i(x)^{\epsilon-1} - f_{ji}.$$

Similar to the domestic profit, the cutoff human capital of the marginal exporter is a sufficient statistics for the size of the foreign market and the marginal cost of accessing to that market. To see this, we start with the cutoff condition:

$$\frac{H_j}{\epsilon}\left(\frac{P_j}{\tau_{ji}w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1}A_i(x^e_{ji})^{\epsilon-1}-f_{ji}w_i=0,$$
$$\frac{H_j}{w_i\epsilon}\left(\frac{P_j}{\tau_{ji}w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1}=\frac{f_{ji}}{b_i e^{(\epsilon-1)(x-x^e_{ji})}}.$$

Plugging the above equation into the original profit-to-wage ratio, we have:

$$\frac{\pi^e_{ji}(x)}{w_i}=f_{ji}[e^{(\epsilon-1)(x-x^e_{ji})}-1].$$

This ratio depends positively on x and negatively on x^e_{ji}. x^e_{ji} is a measure of the access to the foreign market: it will be lower (easier to access) when τ_{ji} is lower, or the foreign market is larger (H_j or P_j higher). When τ_{ji} is lower, the profit-to-wage ratio from the exporting market will be higher.

FDI Profits The profits earned from FDI to the foreign country, divided by local wage, is:

$$\frac{\pi^f_{ji}(x)}{w_i}=\frac{H_j}{w_i\epsilon}\left(\frac{P_j}{w_j}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1}A_i(x)^{\epsilon-1}-g_{ji}.$$

From the FDI cutoff condition, we know

$$\frac{H_j}{w_i\epsilon}\left(\frac{P_j}{w_j}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1}A_i(x^f_{ji})^{\epsilon-1}=\frac{H_j}{w_i\epsilon}\left(\frac{P_j}{\tau_{ji}w_i}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1}A_i(x^f_{ji})^{\epsilon-1}+(g_{ji}-f_{ji}),$$
$$\frac{H_j}{w_i\epsilon}\left(\frac{P_j}{w_j}\frac{\epsilon-1}{\epsilon}\right)^{\epsilon-1}=\left[f_{ji}\frac{A_i(x^f_{ji})^{\epsilon-1}}{A_i(x^e_{ji})^{\epsilon-1}}+g_{ji}-f_{ji}\right]\frac{1}{A_i(x^f_{ji})^{\epsilon-1}}.$$

Therefore

$$\frac{\pi^f_{ji}(x)}{w_i}=f_{ji}e^{(\epsilon-1)(x-x^e_{ji})}+(g_{ji}-f_{ji})e^{(\epsilon-1)(x-x^f_{ji})}-g_{ji}.$$

This profit-to-wage ratio decreases with x_{ji}^f:

$$\frac{\partial \frac{\pi_{ji}^f(x)}{w_i}}{\partial x_{ji}^f} = e^{(\epsilon-1)(x-x_{ji}^f)}(g_{ji}-f_{ji})(1-\epsilon) < 0.$$

A.2 Calibration

The TFP b_i and measure of population are computed following the method in *Caselli* (2005). The computation is based on Penn World Table 7.0, and all undefined variable names in italics are the standard variable names in PWT. I first compute real GDP in year t, Y_t, as

$$Y_t = pop_t \cdot rgdpl_t.$$

The number of workers, L_t, is backed out by

$$L_t = Y_t / rgdpwok_t.$$

This raw measure of the stock of work-force is first adjusted by human capital. Using years of school attainment for both males and females 25 years old and above from *Barro and Lee* (2010), I construct human capital h_t as

$$h_t = e^{\phi(c_t)},$$

where c_t is the years of schooling and $\phi(c_t)$ is piece-wise linear:

$$\phi(c_t) = \begin{cases} 0.134 * c & \text{if} \quad c_t \leq 4 \\ 0.134 * 4 + 0.101 * (c_t - 4) & \text{if} \quad 4 < c_t \leq 8 \\ 0.134 * 4 + 0.101 * 4 * 0.068 * (c_t - 4) & \text{if} \quad 8 < c_t \end{cases} \cdot$$

Because the year of schooling data are only available at five-year intervals, linear interpolation is used to fill in the gap years. c_t is a slow-moving variable; therefore, linear interpolation can provide reasonably smooth estimations.

To construct the stock of physical capital in each year, I first compute investment in each year as

$$I_t = Y_t * \mathrm{ki_t}/100,$$

and then back out the initial capital stock using perpetual inventory method. I assume that capital and output grow at the same rate, and the depreciation rate is 6 percent per year. The initial capital stock when $t = 0$ is

$$K_0 = I_0/(g_k + 0.06),$$

where g_k is the average growth rate of GDP in the first 10 years of data. Given the initial capital stock, the sequence of capital stock in year t is computed as

$$K_t = (1 - 0.06)K_{t-1} + I_t.$$

With a computed sequence of physical capital, the final measure of population year t, n_t, is computed as

$$n_t = K_t^a (h_t L_t)^{1-a},$$

where $a = 1/3$ and the TFP, b_t, is calculated as

$$b_t = Y_t/n_t.$$

At the end, b_t is normalized so that the TFP for the U.S. in 1988 is 1. For the sequence of estimated TFP, see Table A.16. Given a sequence of n_t for each country, I first average across the years to get a single measure for each country. I then normalize across the countries so n_{USA} is 1.

A.3 Tables and Figures

A.3.1 Empirical Results: Public Firms

	Matched Data		ExecuCompustat	
Sector	Percent	N.Obs.	Percent	N.Obs.
Mineral & Construction	4.39%	751	5.44%	1876
Manufacturing	46.15%	7892	42.51%	14649
Transportation, Communications and Utilities	10.79%	1845	11.24%	3873
Wholesale and Retail Trade	12.36%	2113	11.49%	3960
Finance, Insurance and Real Estate	13.91%	2379	15.28%	5265
Services	12.40%	2121	14.03%	4835
Other	0.71%	122	0.69%	239
Total	100.00%	17223	100.00%	34697

Table A.1: Sector Composition: Public Firm Sample

Note: This table reports the sectoral composition of the firm-year observations in the linked ExecuCompustat-LBD-LFTTD data set and compares the distribution with the original ExecuCompustat data set. The sector definition is based on a one-digit SIC code.

Mean	Exporters	Non-Exporters	Overall
CEO Compensation, Estimated	4487.7	3254.3	4197.1
CEO Compensation, Realized	4662.4	3340.4	4350.8
CEO-to-worker Pay Ratio, Estimated	91.9	80.8	89.3
CEO-to-worker Pay Ratio, Realized	91.8	79.6	88.9
N. Observations	13169	4054	17223

Table A.2: Summary Statistics: Public Firm Sample

Note: This table reports the mean of key variables of the linked ExecuCompustat-LBD-LFTTD data set. The unit of observation is firm-year. Executive compensations are measured in thousands of U.S. dollars. For the difference between estimated and realized compensation, see Section 2.2.

(a) Estimated Compensation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.486***	2.22e-05		0.0265		0.0168		0.0478*		0.00693	
	(0.0292)	(0.0226)		(0.0226)		(0.0245)		(0.0255)		(0.0214)	
Sales		0.438***	0.438***							0.156***	0.156***
		(0.00712)	(0.00704)							(0.0215)	(0.0215)
Asset				0.434***	0.436***					0.302***	0.302***
				(0.00660)	(0.00654)					(0.0204)	(0.0204)
Employment						0.398***	0.399***			0.654***	0.654***
						(0.00706)	(0.00694)			(0.0264)	(0.0265)
Payroll								0.365***	0.369***	-0.654***	-0.654***
								(0.00834)	(0.00810)	(0.0292)	(0.0292)
Constant	2.066***	-0.107	-0.107	-0.0754	-0.0614	-0.458**	-0.452**	-1.533***	-1.529***	2.100***	2.101***
	(0.194)	(0.184)	(0.183)	(0.185)	(0.185)	(0.187)	(0.186)	(0.199)	(0.200)	(0.222)	(0.222)
Observations	17223	17223	17223	17223	17223	17223	17223	17223	17223	17223	17223
R-squared	0.266	0.459	0.459	0.457	0.457	0.429	0.429	0.397	0.397	0.502	0.502

(b) Realized Compensation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.507***	0.0238		0.0621**		0.0496*		0.0700***		0.0273	
	(0.0300)	(0.0246)		(0.0247)		(0.0257)		(0.0270)		(0.0234)	
Sales		0.436***	0.437***							0.249***	0.250***
		(0.00719)	(0.00702)							(0.0209)	(0.0209)
Asset				0.420***	0.425***					0.191***	0.191***
				(0.00692)	(0.00676)					(0.0201)	(0.0201)
Employment						0.388***	0.391***			0.526***	0.526***
						(0.00731)	(0.00706)			(0.0320)	(0.0320)
Payroll								0.364***	0.370***	-0.516***	-0.514***
								(0.00864)	(0.00828)	(0.0356)	(0.0355)
Constant	2.017***	-0.144	-0.132	-0.0576	-0.0248	-0.444**	-0.425**	-1.574***	-1.569***	1.586***	1.590***
	(0.210)	(0.196)	(0.196)	(0.202)	(0.201)	(0.203)	(0.203)	(0.216)	(0.216)	(0.245)	(0.245)
Observations	17223	17223	17223	17223	17223	17223	17223	17223	17223	17223	17223
R-squared	0.270	0.439	0.439	0.428	0.428	0.407	0.407	0.385	0.385	0.461	0.461

Table A.3: CEO-to-Worker Pay Ratio: U.S. Public Firms by Exporting Status

Note: This table reports the results of estimating equation (2.1) for U.S. public firms based on the linked ExecuCompustat-LBD-LFTTD data. The LHS variable for each of the regressions is the (log of) CEO-to-worker pay ratio. The upper panel uses estimated compensation on the LHS, and the lower panel uses realized compensation on the LHS. For the difference between the two, refer to Section 2.2. "Exporter" is the exporter indicator computed from LFTTD. "Sales" is the (log of) total annual sales reported in ExecuCompustat. "Asset" is the (log of) total asset reported in ExecuCompustat. "Employment" is the (log of) March 12 employment reported in LBD at the firm level. "Payroll" is the (log of) total annual payroll reported in LBD. The unit of observation is firm-year. All regressions include year and four-digit SIC fixed effects. Robust standard errors are clustered at the year-sector level.

(a) Estimated Compensation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.459***	0.00504		0.0202		0.0428**		0.0760***		0.0251	
	(0.0263)	(0.0195)		(0.0193)		(0.0216)		(0.0225)		(0.0177)	
Sales		0.414***	0.414***							0.146***	0.147***
		(0.00636)	(0.00628)							(0.0193)	(0.0193)
Asset				0.414***	0.416***					0.345***	0.345***
				(0.00498)	(0.00495)					(0.0166)	(0.0166)
Employment						0.356***	0.359***			0.623***	0.623***
						(0.00665)	(0.00663)			(0.0236)	(0.0237)
Payroll								0.323***	0.329***	-0.691***	-0.689***
								(0.00788)	(0.00772)	(0.0272)	(0.0272)
Constant	2.083***	-0.00830	-0.00563	0.0258	0.0369	-0.203*	-0.185	-1.116***	-1.108***	2.482***	2.486***
	(0.130)	(0.120)	(0.120)	(0.118)	(0.117)	(0.117)	(0.116)	(0.139)	(0.139)	(0.159)	(0.158)
Observations	16268	16268	16268	16268	16268	16268	16268	16268	16268	16268	16268
R-squared	0.356	0.602	0.602	0.605	0.605	0.543	0.543	0.502	0.502	0.664	0.664

(b) Realized Compensation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.477***	0.00622		0.0336*		0.0585***		0.0743***		0.0221	
	(0.0268)	(0.0202)		(0.0201)		(0.0221)		(0.0232)		(0.0188)	
Sales		0.430***	0.430***							0.250***	0.250***
		(0.00624)	(0.00609)							(0.0197)	(0.0197)
Asset				0.419***	0.422***					0.238***	0.238***
				(0.00553)	(0.00540)					(0.0173)	(0.0173)
Employment						0.358***	0.363***			0.444***	0.444***
						(0.00688)	(0.00679)			(0.0290)	(0.0290)
Payroll								0.340***	0.345***	-0.497***	-0.495***
								(0.00841)	(0.00813)	(0.0333)	(0.0331)
Constant	2.001***	-0.170	-0.167	-0.0800	-0.0614	-0.300**	-0.276**	-1.367***	-1.359***	1.633***	1.637***
	(0.144)	(0.129)	(0.128)	(0.130)	(0.130)	(0.135)	(0.134)	(0.155)	(0.155)	(0.183)	(0.184)
Observations	16268	16268	16268	16268	16268	16268	16268	16268	16268	16268	16268
R-squared	0.341	0.575	0.575	0.566	0.566	0.508	0.508	0.484	0.483	0.602	0.602

Table A.4: Top-Five-Executives-to-Worker Pay Ratio: U.S. Public Firms by Exporting Status

Note: This table reports the results of estimating equation (2.1) based on the ExecuCompustat-LBD-LFTTD data. The LHS variable is the (log of) average compensation of the top five highly paid executives divided by the average wage. For other details, see the note to Table A.3.

(a) Salary

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.213***	-0.0340*		0.00479		-0.0822***		0.0104		-0.0161	
	(0.0227)	(0.0199)		(0.0206)		(0.0201)		(0.0211)		(0.0153)	
Sales		0.223***	0.220***							0.155***	0.155***
		(0.00619)	(0.00636)							(0.0117)	(0.0118)
Asset				0.197***	0.197***					0.0699***	0.0701***
				(0.00644)	(0.00660)					(0.0109)	(0.0109)
Employment						0.250***	0.244***			1.058***	1.058***
						(0.00623)	(0.00619)			(0.0268)	(0.0268)
Payroll								0.169***	0.170***	-1.055***	-1.056***
								(0.00744)	(0.00745)	(0.0271)	(0.0273)
Constant	1.767***	0.669***	0.652***	0.804***	0.807***	0.188	0.156	0.110	0.111	4.342***	4.340***
	(0.135)	(0.131)	(0.130)	(0.136)	(0.134)	(0.133)	(0.132)	(0.151)	(0.151)	(0.173)	(0.173)
Observations	17156	17156	17156	17156	17156	17156	17156	17156	17156	17156	17156
R-squared	0.370	0.438	0.438	0.423	0.423	0.458	0.457	0.408	0.408	0.570	0.570

(b) Bonus

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.508***	-0.00229		0.0503*		0.0294		0.0759**		0.00726	
	(0.0352)	(0.0278)		(0.0283)		(0.0301)		(0.0316)		(0.0266)	
Sales		0.454***	0.454***							0.318***	0.318***
		(0.00797)	(0.00770)							(0.0254)	(0.0255)
Asset				0.429***	0.433***					0.163***	0.162***
				(0.00785)	(0.00760)					(0.0231)	(0.0231)
Employment						0.411***	0.414***			0.713***	0.713***
						(0.00838)	(0.00817)			(0.0343)	(0.0343)
Payroll								0.370***	0.376***	-0.736***	-0.735***
								(0.00974)	(0.00944)	(0.0362)	(0.0362)
Constant	1.078***	-0.963***	-0.964***	-0.848***	-0.819***	-1.416***	-1.403***	-2.482***	-2.471***	1.679***	1.681***
	(0.198)	(0.179)	(0.178)	(0.184)	(0.183)	(0.179)	(0.178)	(0.202)	(0.202)	(0.223)	(0.223)
Observations	12681	12681	12681	12681	12681	12681	12681	12681	12681	12681	12681
R-squared	0.340	0.502	0.502	0.485	0.485	0.476	0.476	0.444	0.444	0.538	0.538

Table A.5: Decomposition of CEO Compensation: Salary and Bonus in U.S. Public Firms

Note: This table reports the results of estimating equation (2.1) based on the ExecuCompustat-LBD-LFTTD data. The LHS variable for the upper panel is the (log of) annual salary of the CEO divided by average wage. The LHS variable for the lower panel is the (log of) annual bonus of the CEO divided by average wage. For other details, see the note to Table A.3.

(a) Estimated

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.676***	-0.00181		0.0119		0.0661		0.0569		-0.000641	
	(0.0546)	(0.0494)		(0.0490)		(0.0522)		(0.0525)		(0.0495)	
Sales		0.612***	0.612***							0.139***	0.139***
		(0.0130)	(0.0127)							(0.0414)	(0.0415)
Asset				0.626***	0.627***					0.515***	0.515***
				(0.0122)	(0.0118)					(0.0399)	(0.0400)
Employment						0.519***	0.524***			0.420***	0.420***
						(0.0130)	(0.0126)			(0.0489)	(0.0489)
Payroll								0.518***	0.522***	-0.433***	-0.433***
								(0.0140)	(0.0134)	(0.0534)	(0.0534)
Constant	-0.799*	-3.859***	-3.860***	-3.891***	-3.885***	-4.114***	-4.087***	-5.922***	-5.917***	-2.439***	-2.439***
	(0.421)	(0.394)	(0.393)	(0.389)	(0.388)	(0.405)	(0.405)	(0.420)	(0.420)	(0.444)	(0.444)
Observations	16963	16963	16963	16963	16963	16963	16963	16963	16963	16963	16963
R-squared	0.183	0.302	0.302	0.309	0.309	0.271	0.271	0.267	0.266	0.315	0.315

(b) Realized

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.855***	0.0232		0.0733		0.119*		0.0682		-0.00242	
	(0.0711)	(0.0678)		(0.0676)		(0.0706)		(0.0713)		(0.0685)	
Sales		0.751***	0.753***							0.417***	0.417***
		(0.0186)	(0.0179)							(0.0629)	(0.0630)
Asset				0.736***	0.741***					0.293***	0.293***
				(0.0177)	(0.0169)					(0.0582)	(0.0582)
Employment						0.626***	0.634***			0.0867	0.0867
						(0.0186)	(0.0177)			(0.0713)	(0.0713)
Payroll								0.657***	0.663***	-0.0146	-0.0148
								(0.0196)	(0.0184)	(0.0780)	(0.0776)
Constant	-0.721	-4.474***	-4.462***	-4.359***	-4.319***	-4.718***	-4.670***	-7.229***	-7.222***	-4.662***	-4.662***
	(0.467)	(0.453)	(0.452)	(0.449)	(0.448)	(0.460)	(0.459)	(0.482)	(0.482)	(0.544)	(0.544)
Observations	16963	16963	16963	16963	16963	16963	16963	16963	16963	16963	16963
R-squared	0.182	0.275	0.275	0.272	0.272	0.248	0.248	0.252	0.252	0.277	0.277

Table A.6: Decomposition of CEO Compensation: Stock and Option Rewards in U.S. Public Firms

Note: This table reports the results of estimating equation (2.1) based on the ExecuCompustat-LBD-LFTTD data. The LHS variable for the upper panel is the (log of) estimated income from stocks and options of the CEO divided by average wage. The LHS variable for the lower panel is the (log of) realized income from stocks and options of the CEO divided by average wage. For other details, see the note to Table A.3.

(a) Estimated

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
MNE	0.284***	0.0192		0.00905		0.140***		0.139***		0.0256	
	(0.0267)	(0.0221)		(0.0228)		(0.0229)		(0.0237)		(0.0217)	
Sales		0.446***	0.447***							0.186***	0.187***
		(0.00814)	(0.00802)							(0.0255)	(0.0254)
Asset				0.433***	0.434***					0.284***	0.285***
				(0.00786)	(0.00775)					(0.0242)	(0.0242)
Employment						0.391***	0.395***			0.696***	0.695***
						(0.00821)	(0.00817)			(0.0311)	(0.0311)
Payroll								0.355***	0.360***	-0.709***	-0.709***
								(0.0100)	(0.00996)	(0.0342)	(0.0342)
Constant	2.274***	-0.477**	-0.478**	-0.349	-0.349	-0.616***	-0.609***	-1.612***	-1.626***	2.013***	2.013***
	(0.273)	(0.231)	(0.231)	(0.236)	(0.236)	(0.232)	(0.234)	(0.249)	(0.251)	(0.278)	(0.278)
Observations	12943	12943	12943	12943	12943	12943	12943	12943	12943	12943	12943
R-squared	0.279	0.473	0.473	0.466	0.466	0.440	0.439	0.406	0.404	0.517	0.517

(b) Realized

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
MNE	0.197***	-0.0669***		-0.0709***		0.0562**		0.0515**		-0.0576**	
	(0.0285)	(0.0245)		(0.0254)		(0.0252)		(0.0259)		(0.0242)	
Sales		0.444***	0.441***							0.277***	0.275***
		(0.00831)	(0.00816)							(0.0253)	(0.0253)
Asset				0.422***	0.418***					0.175***	0.172***
				(0.00821)	(0.00807)					(0.0247)	(0.0246)
Employment						0.382***	0.383***			0.543***	0.545***
						(0.00852)	(0.00844)			(0.0383)	(0.0383)
Payroll								0.357***	0.359***	-0.542***	-0.542***
								(0.0103)	(0.0102)	(0.0421)	(0.0421)
Constant	2.338***	-0.400	-0.397	-0.215	-0.211	-0.483*	-0.480*	-1.564***	-1.569***	1.491***	1.490***
	(0.288)	(0.252)	(0.251)	(0.261)	(0.261)	(0.262)	(0.262)	(0.279)	(0.279)	(0.304)	(0.304)
Observations	12943	12943	12943	12943	12943	12943	12943	12943	12943	12943	12943
R-squared	0.277	0.443	0.443	0.431	0.430	0.410	0.410	0.387	0.387	0.466	0.465

Table A.7: CEO-to-Worker Pay Ratio: U.S. Public Firms by Multinational Status

A.3.2 Empirical Results: Private Firms

	Matched Data		Capital IQ	
Sector	Percent	N.Obs.	Percent	N.obs.
Mineral & Construction	3.32%	199	4.13%	483
Manufacturing	33.86%	2032	34.44%	4032
Transportation, Communications and Utilities	10.71%	643	10.23%	1197
Wholesale and Retail Trade	9.30%	558	9.18%	1075
Finance, Insurance and Real Estate	21.98%	1319	18.85%	2206
Services	19.99%	1200	21.80%	2552
Other	0.85%	51	1.38%	161
Total	100.00%	6002	100.00%	11706

Table A.8: Sector Composition: Private Firm Sample

Note: This table reports the sectoral composition of the firm-year observations in the linked CIQ-LBD-LFTTD data set and compares the distribution with the original Capital-IQ data set. The sector definition is based on one-digit SIC code.

Mean	Exporters	Non-Exporters	Overall
Top 1 Compensation, Estimated	2626.9	1731.2	2233.5
Top 1 Compensation, Realized	2157	1522.1	1878.2
Top-1-to-worker Pay Ratio, Estimated	49.8	36.7	44
Top-1-to-worker Pay Ratio, Realized	41.3	32.8	37.6
N. Observations	3366	2636	6002

Table A.9: Summary Statistics: Private Firm Sample

Note: This table reports the mean of key variables of the linked CIQ-LBD-LFTTD data set. The unit of observation is firm-year. Executive compensations are measured in thousands of U.S. dollars. For the difference between estimated and realized compensation, see Section 2.2.

(a) Estimated Compensation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.442***	-0.0683*		-0.0398		-0.0834**		-0.0919**		-0.0933***	
	(0.0520)	(0.0392)		(0.0401)		(0.0399)		(0.0420)		(0.0340)	
Sales		0.417***	0.413***							0.166***	0.163***
		(0.0104)	(0.0101)							(0.0215)	(0.0216)
Asset				0.414***	0.412***					0.228***	0.230***
				(0.00895)	(0.00879)					(0.0185)	(0.0185)
Employment						0.384***	0.378***			0.627***	0.632***
						(0.0101)	(0.00951)			(0.0362)	(0.0365)
Payroll								0.373***	0.366***	-0.552***	-0.564***
								(0.0112)	(0.0105)	(0.0407)	(0.0406)
Constant	2.810***	0.824***	0.810***	0.801***	0.792***	0.316	0.312	-3.477***	-3.399***	6.148***	6.272***
	(0.160)	(0.193)	(0.190)	(0.174)	(0.172)	(0.201)	(0.198)	(0.275)	(0.269)	(0.473)	(0.472)
Observations	6002	6002	6002	6002	6002	6002	6002	6002	6002	6002	6002
R-squared	0.363	0.595	0.595	0.596	0.596	0.559	0.558	0.533	0.532	0.651	0.651

(b) Realized Compensation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.411***	-0.0630		-0.0441		-0.0979**		-0.0932**		-0.0803**	
	(0.0546)	(0.0415)		(0.0419)		(0.0397)		(0.0426)		(0.0325)	
Sales		0.388***	0.384***							0.120***	0.118***
		(0.0108)	(0.0105)							(0.0210)	(0.0211)
Asset				0.391***	0.389***					0.249***	0.251***
				(0.00910)	(0.00916)					(0.0180)	(0.0180)
Employment						0.372***	0.365***			0.731***	0.735***
						(0.00988)	(0.00917)			(0.0328)	(0.0333)
Payroll								0.353***	0.345***	-0.661***	-0.671***
								(0.0112)	(0.0105)	(0.0385)	(0.0389)
Constant	2.690***	0.844***	0.831***	0.792***	0.782***	0.273	0.269	-3.248***	-3.169***	7.298***	7.404***
	(0.199)	(0.233)	(0.230)	(0.221)	(0.219)	(0.243)	(0.239)	(0.307)	(0.300)	(0.480)	(0.482)
Observations	6002	6002	6002	6002	6002	6002	6002	6002	6002	6002	6002
R-squared	0.402	0.619	0.618	0.627	0.627	0.601	0.600	0.566	0.565	0.696	0.695

Table A.10: Top-1-to-Worker Pay Ratio: U.S. Private Firms

Note: This table reports the results of estimating equation (2.1) for U.S. private firms based on the linked CIQ-LBD-LFTTD data. The LHS variable for each of the regressions is the (log of) highest total compensation divided by the average income for a given firm within a given year. The upper panel uses estimated compensation on the LHS, and the lower panel uses realized compensation on the LHS. For the difference between the two, refer to Section 2.2. “Exporter” is the exporter indicator computed from LFTTD. “Sales” is the (log of) total annual sales reported in CIQ. “Asset” is the (log of) total asset reported in CIQ. “Employment” is the (log of) March 12 employment reported in LBD at the firm level. “Payroll” is the (log of) total annual payroll reported in LBD. The unit of observation is firm-year. In all the regressions, year and four-digit SIC fixed effects are controlled for. Robust standard errors are clustered at the year-sector level.

(a) Estimated Compensation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.409***	-0.0485		-0.00658		-0.0658*		-0.0708*		-0.0649**	
	(0.0507)	(0.0370)		(0.0382)		(0.0361)		(0.0385)		(0.0291)	
Sales		0.406***	0.402***							0.162***	0.160***
		(0.0122)	(0.0115)							(0.0207)	(0.0207)
Asset				0.400***	0.399***					0.211***	0.212***
				(0.00997)	(0.00947)					(0.0197)	(0.0197)
Employment						0.374***	0.369***			0.623***	0.626***
						(0.00998)	(0.00952)			(0.0338)	(0.0341)
Payroll								0.367***	0.361***	-0.551***	-0.559***
								(0.0110)	(0.0105)	(0.0372)	(0.0374)
Constant	2.215***	0.327***	0.311***	0.334***	0.331***	-0.0530	-0.0670	-3.856***	-3.804***	5.795***	5.873***
	(0.0928)	(0.109)	(0.107)	(0.0934)	(0.0925)	(0.144)	(0.137)	(0.220)	(0.216)	(0.424)	(0.425)
Observations	4827	4827	4827	4827	4827	4827	4827	4827	4827	4827	4827
R-squared	0.411	0.647	0.647	0.644	0.644	0.627	0.627	0.595	0.595	0.712	0.712

(b) Realized Compensation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.390***	-0.0444		-0.0136		-0.0778**		-0.0725*		-0.0546*	
	(0.0540)	(0.0392)		(0.0390)		(0.0370)		(0.0397)		(0.0284)	
Sales		0.385***	0.382***							0.116***	0.114***
		(0.0128)	(0.0120)							(0.0198)	(0.0199)
Asset				0.388***	0.387***					0.244***	0.245***
				(0.00995)	(0.00954)					(0.0184)	(0.0184)
Employment						0.368***	0.362***			0.713***	0.716***
						(0.00976)	(0.00925)			(0.0315)	(0.0320)
Payroll								0.353***	0.347***	-0.648***	-0.654***
								(0.0110)	(0.0106)	(0.0365)	(0.0369)
Constant	1.988***	0.198	0.183	0.164	0.159	-0.244	-0.260	-3.857***	-3.804***	6.681***	6.747***
	(0.122)	(0.160)	(0.158)	(0.142)	(0.141)	(0.171)	(0.164)	(0.244)	(0.238)	(0.438)	(0.440)
Observations	4827	4827	4827	4827	4827	4827	4827	4827	4827	4827	4827
R-squared	0.429	0.657	0.657	0.665	0.665	0.654	0.654	0.613	0.612	0.747	0.746

Table A.11: Top-Five-Executives-to-Worker Pay Ratio: U.S. Private Firms

Note: This table reports the results of estimating equation (2.1) based on the CIQ-LBD-LFTTD data. The LHS variable is the (log of) average compensation of the top five highly paid executives divided by the average wage. For other details, see the note to Table A.10.

(a) Salary

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.168***	-0.0451		-0.0236		-0.135***		-0.0742**		-0.0168	
	(0.0364)	(0.0349)		(0.0349)		(0.0305)		(0.0341)		(0.0229)	
Sales		0.176***	0.173***							0.0527***	0.0522***
		(0.00843)	(0.00779)							(0.0162)	(0.0162)
Asset				0.164***	0.162***					0.105***	0.105***
				(0.00908)	(0.00872)					(0.0151)	(0.0151)
Employment						0.223***	0.213***			0.954***	0.955***
						(0.00740)	(0.00698)			(0.0219)	(0.0216)
Payroll								0.169***	0.163***	-0.908***	-0.910***
								(0.00841)	(0.00773)	(0.0257)	(0.0250)
Constant	2.070***	1.237***	1.228***	1.282***	1.276***	0.628***	0.621***	-0.778***	-0.716***	10.40***	10.43***
	(0.0430)	(0.0672)	(0.0662)	(0.0702)	(0.0696)	(0.0843)	(0.0795)	(0.155)	(0.148)	(0.292)	(0.285)
Observations	5123	5123	5123	5123	5123	5123	5123	5123	5123	5123	5123
R-squared	0.497	0.573	0.573	0.566	0.566	0.620	0.618	0.563	0.562	0.738	0.738

(b) Bonus

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.313***	-0.237***		-0.214***		-0.260***		-0.270***		-0.243***	
	(0.0855)	(0.0648)		(0.0634)		(0.0655)		(0.0693)		(0.0613)	
Sales		0.490***	0.474***							0.146***	0.138***
		(0.0227)	(0.0216)							(0.0383)	(0.0378)
Asset				0.482***	0.469***					0.302***	0.307***
				(0.0210)	(0.0206)					(0.0337)	(0.0334)
Employment						0.441***	0.423***			0.775***	0.791***
						(0.0213)	(0.0207)			(0.0629)	(0.0632)
Payroll								0.430***	0.408***	-0.695***	-0.728***
								(0.0245)	(0.0233)	(0.0771)	(0.0758)
Constant	1.587***	-0.766***	-0.809***	-0.768***	-0.813***	-1.292***	-1.306***	-5.660***	-5.428***	6.085***	6.436***
	(0.0702)	(0.137)	(0.134)	(0.141)	(0.137)	(0.175)	(0.170)	(0.429)	(0.415)	(0.915)	(0.900)
Observations	3927	3927	3927	3927	3927	3927	3927	3927	3927	3927	3927
R-squared	0.388	0.546	0.544	0.551	0.549	0.534	0.531	0.511	0.508	0.595	0.593

Table A.12: Decomposition of the Highest Compensation: Salary and Bonus in U.S. Private Firms

Note: This table reports the results of estimating equation (2.1) based on the CIQ-LBD-LFTTD data. The LHS variable for the upper panel is the (log of) annual salary of the highest-paid executive divided by average wage. The LHS variable for the lower panel is the (log of) annual bonus of the highest-paid executive divided by average wage. For other details, see the note to Table A.10.

(a) Estimated

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.608***	-0.106		-0.0605		-0.0957		-0.165*		-0.176**	
	(0.104)	(0.0909)		(0.0914)		(0.0905)		(0.0925)		(0.0888)	
Sales		0.603***	0.596***							0.271***	0.265***
		(0.0209)	(0.0198)							(0.0466)	(0.0463)
Asset				0.596***	0.592***					0.273***	0.278***
				(0.0189)	(0.0180)					(0.0444)	(0.0443)
Employment						0.507***	0.501***			0.352***	0.363***
						(0.0191)	(0.0184)			(0.0801)	(0.0799)
Payroll								0.532***	0.518***	-0.228***	-0.253***
								(0.0190)	(0.0181)	(0.0863)	(0.0857)
Constant	1.057***	-1.764***	-1.802***	-1.732***	-1.755***	-2.006***	-2.027***	-7.699***	-7.586***	0.142	0.373
	(0.333)	(0.372)	(0.372)	(0.331)	(0.329)	(0.332)	(0.329)	(0.458)	(0.460)	(0.987)	(0.985)
Observations	4742	4742	4742	4742	4742	4742	4742	4742	4742	4742	4742
R-squared	0.319	0.471	0.470	0.470	0.470	0.431	0.431	0.431	0.431	0.486	0.485

(b) Realized

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Exporter	0.544***	-0.0457		-0.0281		-0.0802		-0.115		-0.0906	
	(0.103)	(0.0836)		(0.0858)		(0.0795)		(0.0817)		(0.0777)	
Sales		0.498***	0.495***							0.125***	0.122***
		(0.0215)	(0.0215)							(0.0427)	(0.0425)
Asset				0.510***	0.509***					0.336***	0.338***
				(0.0201)	(0.0206)					(0.0442)	(0.0441)
Employment						0.450***	0.445***			0.570***	0.576***
						(0.0200)	(0.0195)			(0.0708)	(0.0710)
Payroll								0.454***	0.444***	-0.468***	-0.480***
								(0.0205)	(0.0201)	(0.0785)	(0.0786)
Constant	0.401	-1.929***	-1.945***	-1.987***	-1.998***	-2.317***	-2.335***	-7.068***	-6.989***	2.507**	2.626**
	(0.699)	(0.742)	(0.740)	(0.723)	(0.721)	(0.723)	(0.721)	(0.801)	(0.801)	(1.121)	(1.124)
Observations	4742	4742	4742	4742	4742	4742	4742	4742	4742	4742	4742
R-squared	0.393	0.500	0.500	0.507	0.507	0.484	0.484	0.477	0.477	0.523	0.523

Table A.13: Decomposition of the Highest Compensation: Stock and Option Rewards in U.S. Private Firms

Note: This table reports the results of estimating equation (2.1) based on the CIQ-LBD-LFTTD data. The LHS variable for the upper panel is the (log of) estimated income from stocks and options of the highest-paid executive divided by average wage. The LHS variable for the lower panel is the (log of) realized income from stocks and options of the highest-paid executive divided by average wage. For other details, see the note to Table A.10.

A.3.3 Detailed Tables and Figures for the Model

Parameter	Value	Target/Source
λ	3.81	Firms size distribution estimated from LBD
ϵ	4.0	Average mark-up
α	28.0	Corporate sales as a percentage of all firms sales
β	0.7373	Estimated from ExecuCompustat
f_{11}	6.0	World Bank Doing Business Index
f_{12}	19.6	World Bank Doing Business Index
f_{21}	24.6	World Bank Doing Business Index
f_{22}	38.9	World Bank Doing Business Index
f-Scale	5.5	Exporter and MNE employment share
n_{ROW}	6.0	*Caselli* (2005),*Barro and Lee* (2010)
n_{USA}	1.0	*Caselli* (2005),*Barro and Lee* (2010)
b_{ROW}	0.57	*Caselli* (2005),*Barro and Lee* (2010)
b_{USA}	1.0	*Caselli* (2005),*Barro and Lee* (2010)
s	2.8	Highest-CEO-to-average-wage ratio among public firms

Table A.14: Calibration Targets and Results

Note: λ is the shape parameter of the exponential distribution. ϵ is the elasticity of substitution in the utility functions. α is the size of the smallest public firm. β is the elasticity of CEO income with respect to firm profit. f_{ij} is the fixed cost of exporting from country j to country i. f-Scale is the normalizing factor of the entire f_{ij} matrix. I divide the f_{ij} matrix by this number. n_i is the measure of capital-adjusted endowment of human capital in country i. b_i is the TFP in country i. s is the upper bound of human capital distribution. See Section 2.5 and Appendix A.2 for the details of calibration. See Table A.16 for the calibrated values of τ, g and TFP by year.

Country	GDP(Bil.$)	Pop(Mil.)	Country	GDP(Bil.$)	Pop(Mil.)
Afghanistan	2.473e+01	2.766e+01	Malawi	8.608e+00	1.462e+01
Albania	1.971e+01	2.984e+00	Malaysia	3.199e+02	2.736e+01
Algeria	2.086e+02	3.377e+01	Maldives	1.784e+00	3.859e-01
Argentina	4.647e+02	4.048e+01	Mali	1.257e+01	1.310e+01
Australia	8.484e+02	2.101e+01	Mauritania	5.948e+00	3.055e+00
Austria	3.227e+02	8.206e+00	Mauritius	1.234e+01	1.274e+00
Bangladesh	1.863e+02	1.513e+02	Mexico	1.362e+03	1.100e+02
Belgium	3.763e+02	1.040e+01	Mongolia	1.025e+01	2.996e+00
Benin	1.009e+01	8.533e+00	Morocco	1.068e+02	3.094e+01
Bolivia	3.469e+01	9.601e+00	Mozambique	1.559e+01	2.144e+01
Botswana	2.118e+01	1.952e+00	Namibia	1.032e+01	2.089e+00
Brazil	1.588e+03	1.963e+02	Nepal	3.027e+01	2.820e+01
Bulgaria	7.829e+01	7.263e+00	Netherlands	6.541e+02	1.665e+01
Burundi	3.641e+00	9.139e+00	New Zealand	1.152e+02	4.173e+00
Cameroon	3.245e+01	1.847e+01	Nicaragua	1.251e+01	5.476e+00
Canada	1.254e+03	3.321e+01	Niger	7.881e+00	1.475e+01
Central African	2.493e+00	4.641e+00	Norway	2.395e+02	4.644e+00
Chile	2.045e+02	1.645e+01	Pakistan	3.943e+02	1.785e+02
China	7.920e+03	1.317e+03	Panama	3.372e+01	3.310e+00
Colombia	3.173e+02	4.314e+01	Papua New Guinea	1.470e+01	5.816e+00
Congo	8.111e+00	3.905e+00	Paraguay	2.338e+01	6.203e+00
Costa Rica	5.038e+01	4.393e+00	Peru	1.965e+02	2.835e+01
Cote d'Ivoire	2.618e+01	2.018e+01	Philippines	2.937e+02	9.606e+01
Denmark	1.967e+02	5.485e+00	Poland	6.074e+02	3.850e+01
Dominican	9.314e+01	9.558e+00	Portugal	2.185e+02	1.068e+01
Ecuador	8.718e+01	1.435e+01	Romania	2.243e+02	2.206e+01
Egypt	3.522e+02	7.727e+01	Rwanda	9.865e+00	1.044e+01
El Salvador	3.838e+01	6.006e+00	Saudi Arabia	5.147e+02	2.492e+01
Fiji	3.732e+00	8.605e-01	Senegal	1.699e+01	1.170e+01
Finland	1.864e+02	5.245e+00	Sierra Leone	4.168e+00	5.023e+00
France	2.061e+03	6.406e+01	Singapore	2.323e+02	4.608e+00
Germany	2.838e+03	8.207e+01	Slovak	1.058e+02	5.455e+00
Ghana	4.465e+01	2.343e+01	South Africa	3.708e+02	4.878e+01
Greece	2.941e+02	1.072e+01	Spain	1.338e+03	4.591e+01
Guatemala	7.899e+01	1.300e+01	Sri Lanka	7.776e+01	2.070e+01
Guyana	3.118e+00	7.581e-01	Sudan	8.558e+01	4.168e+01
Haiti	1.249e+01	9.639e+00	Sweden	3.276e+02	9.045e+00
Honduras	2.838e+01	7.676e+00	Switzerland	3.033e+02	7.582e+00
Hong Kong	2.619e+02	7.019e+00	Syria	8.062e+01	2.132e+01
Hungary	1.742e+02	1.002e+01	Tanzania	4.362e+01	4.021e+01
Iceland	1.315e+01	3.044e-01	Thailand	5.029e+02	6.553e+01
India	3.364e+03	1.141e+03	Togo	4.490e+00	6.220e+00
Indonesia	8.589e+02	2.375e+02	Tonga	8.130e-01	1.049e-01
Iran	7.118e+02	7.503e+01	Trinidad &Tobago	3.973e+01	1.231e+00
Iraq	1.133e+02	2.822e+01	Tunisia	6.190e+01	1.032e+01
Ireland	1.808e+02	4.518e+00	Turkey	7.798e+02	7.579e+01
Israel	1.824e+02	7.112e+00	Uganda	3.427e+01	3.137e+01
Italy	1.797e+03	6.009e+01	United Arab Emirates	2.965e+02	4.621e+00
Jamaica	2.535e+01	2.804e+00	United Kingdom	2.160e+03	6.164e+01
Japan	4.122e+03	1.273e+02	United States	1.301e+04	3.044e+02
Jordan	2.740e+01	6.133e+00	Uruguay	3.485e+01	3.286e+00
Kenya	4.436e+01	3.795e+01	Venezuela	2.686e+02	2.641e+01
Korea	1.223e+03	4.838e+01	Vietnam	2.127e+02	8.756e+01
Laos	1.418e+01	6.145e+00	Zambia	1.827e+01	1.269e+01
Lesotho	2.649e+00	1.915e+00	Zimbabwe	3.129e+00	1.135e+01

Table A.15: Countries Included in Calibration

Note: This table reports the list of countries (110 in total) included in the calibration. All the countries except the U.S. are included in ROW. The GDP and population data are based on Penn World Table 7.0 in the year 2008. GDP is in the unit of constant 2005 international dollar and calculated as the product of *RGDPL* and *POP*.

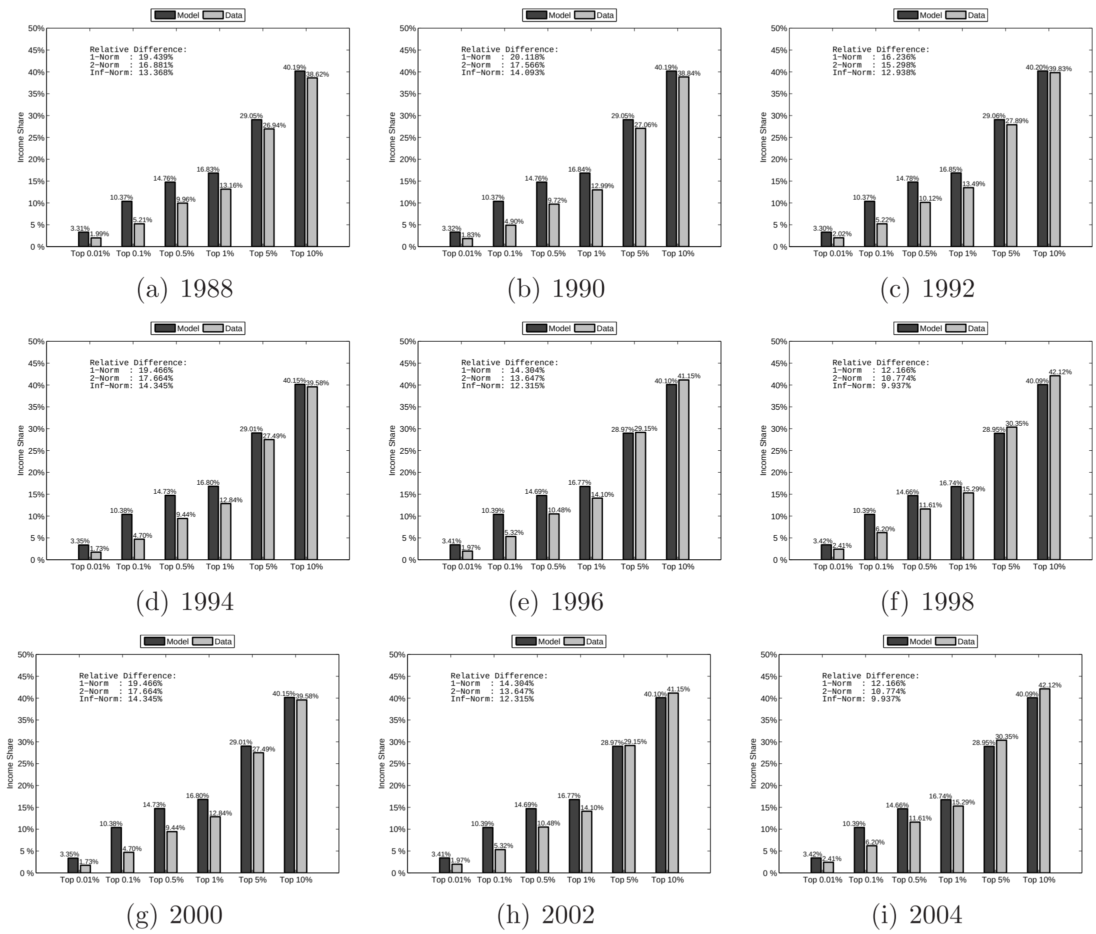

Figure A.1: Model Fit: Top Income Shares (Selected Years)

Note: This figure compares the model-generated top income shares with the data for selected years. For more details see the note to Figure 2.4.

	τ	g	TFP, USA	TFP, ROW
1988	1.835	1523.500	1.000	0.573
1989	1.841	1512.500	1.009	0.570
1990	1.841	1495.700	1.004	0.571
1991	1.863	1488.600	0.986	0.562
1992	1.856	1491.200	1.000	0.559
1993	1.848	1474.900	1.012	0.553
1994	1.822	1451.500	1.031	0.555
1995	1.803	1400.900	1.035	0.554
1996	1.800	1389.000	1.051	0.555
1997	1.787	1376.400	1.069	0.557
1998	1.781	1404.600	1.090	0.551
1999	1.759	1383.100	1.113	0.552
2000	1.714	1354.200	1.119	0.559
2001	1.754	1367.300	1.103	0.555
2002	1.756	1388.400	1.098	0.553
2003	1.748	1342.400	1.101	0.553
2004	1.712	1287.000	1.120	0.560
2005	1.690	1243.200	1.127	0.566
2006	1.673	1217.500	1.131	0.578
2007	1.677	1150.600	1.127	0.591
2008	1.656	1093.300	1.102	0.588

Table A.16: τ and TFP

Note: This table reports the calibrated iceberg trade cost τ and the estimated TFP. The τ and g matrices are assumed to be symmetric. Therefore I only report one off-diagonal term. The calibrated τ and g assume that the TFP for both countries is fixed at the 1988 level. The TFP reported is calculated using the method outlined in *Caselli* (2005) and normalized so that the TFP in the U.S. in 1988 is 1. See Appendix A.2 for details.

BIBLIOGRAPHY

BIBLIOGRAPHY

Anderson, J. E., and E. van Wincoop (2004), Trade costs, *Journal of Economic Literature*, *42*(3), 691–751.

Arnold, B. C. (1985), *Pareto Distributions*, International Co-operative Publishing House.

Axtell, R. L. (2001), Zipf distribution of u.s. firm sizes, *Science*, *293*(5536), 1818–1820.

Bachmann, R., and L. Ma (2012), Lumpy investment, lumpy inventories, *NBER Working Papers 17924*, National Bureau of Economic Research, Inc.

Bachmann, R., E. Engel, and R. Caballero (2013), Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model, *American Economic Journal: Macroeconomics*, *5*(4), 29–67.

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1992), International real business cycles, *Journal of Political Economy*, *100*(4), 745–75.

Baker, G. P., and B. J. Hall (2004), Ceo incentives and firm size, *Journal of Labor Economics*, *22*(4), 767–798.

Barro, R. J., and J.-W. Lee (2010), A new data set of educational attainment in the world, 1950-2010, *NBER Working Papers 15902*, National Bureau of Economic Research, Inc.

Belliveau, M. A., C. A. O'Reilly, and J. B. Wade (1996), Social capital at the top: Effects of social similarity and status on ceo compensation, *The Academy of Management Journal*, *39*(6), pp. 1568–1593.

Bernard, A. B., and J. B. Jensen (1999), Exceptional exporter performance: Cause, effect, or both?, *Journal of International Economics*, *47*(1), 1–25.

Bernard, A. B., J. B. Jensen, and P. K. Schott (2009), Importers, exporters and multinationals: A portrait of firms in the u.s. that trade goods, in *Producer Dynamics: New Evidence from Micro Data*, NBER Chapters, pp. 513–552, National Bureau of Economic Research, Inc.

Bloom, N. (2009), The Impact of Uncertainty Shocks, *Econometrica*, *77*(3), 623–685.

Burstein, A., and J. Vogel (2012), International trade, technology, and the skill premium, *2012 Meeting Papers 664*, Society for Economic Dynamics.

Caballero, R. (2010), Macroeconomics after the crisis: Time to deal with the pretense-of-knowledge syndrome, *Journal of Economic Perspectives*, *24*(4), 85–102.

Caballero, R., and E. Engel (1999), Explaining Investment Dynamics in US Manufacturing: a Generalized (S,s) Approach, *Econometrica*, *67*(4), 783–826.

Caballero, R. J., and E. M. R. A. Engel (1993), Microeconomic adjustment hazards and aggregate dynamics, *The Quarterly Journal of Economics*, *108*(2), 359–83.

Caplin, A. S., and D. F. Spulber (1987), Menu costs and the neutrality of money, *The Quarterly Journal of Economics*, *102*(4), pp. 703–726.

Caselli, F. (2005), Accounting for cross-country income differences, in *Handbook of Economic Growth*, *Handbook of Economic Growth*, vol. 1, edited by P. Aghion and S. Durlauf, chap. 9, pp. 679–741, Elsevier.

Clementi, F., and M. Gallegati (2005), Pareto's law of income distribution: Evidence for grermany, the united kingdom, and the united states, *Microeconomics 0505006*, EconWPA.

Cooper, R., and J. Haltiwanger (2006), On the Nature of Capital Adjustment Costs, *Review of Economic Studies*, *73*(3), 611–633.

di Giovanni, J., A. A. Levchenko, and R. Rancire (2011), Power laws in firm size and openness to trade: Measurement and implications, *Journal of International Economics*, *85*(1), 42–52.

Domowitz, I., R. G. Hubbard, and B. C. Petersen (1988), Market structure and cyclical fluctuations in u.s. manufacturing, *The Review of Economics and Statistics*, *70*(1), 55–66.

Doms, M., and T. Dunne (1998), Capital Adjustment Patterns in Manufacturing Plants, *Review of Economic Dynamics*, *1*(2), 409–429.

Drăgulescu, A., and V. M. Yakovenko (2001a), Evidence for the exponential distribution of income in the usa, *The European Physical Journal B-Condensed Matter and Complex Systems*, *20*(4), 585–589.

Drăgulescu, A., and V. M. Yakovenko (2001b), Exponential and power-law probability distributions of wealth and income in the united kingdom and the united states, *Physica A: Statistical Mechanics and its Applications*, *299*(1), 213–221.

Egger, H., and U. Kreickemeier (2012), Fairness, trade, and inequality, *Journal of International Economics*, *86*(2), 184–196.

Faleye, O., E. Reis, and A. Venkateswaran (2013), The determinants and effects of ceo-employee pay ratios, *Journal of Banking & Finance*, *37*(8), 3258–3272.

Feenstra, R. C., and G. H. Hanson (1996), Globalization, outsourcing, and wage inequality, *American Economic Review*, *86*(2), 240–45.

Feenstra, R. C., and D. E. Weinstein (2010), Globalization, markups, and the u.s. price level, *NBER Working Papers 15749*, National Bureau of Economic Research, Inc.

Feller, W. (1966), *An Introduction to Probability Theory and Its Applications*, vol. 2, John Wiley & Sons, Inc.

Fiori, G. (2012), Lumpy Investment Is Relevant for the Business Cycle!, *Journal of Monetary Economics*, forthcoming.

Frydman, C., and R. E. Saks (2010), Executive compensation: A new view from a long-term perspective, 1936–2005, *Review of Financial Studies*, *23*(5), 2099–2138.

Gabaix, X., and A. Landier (2008), Why has ceo pay increased so much?, *The Quarterly Journal of Economics*, *123*(1), 49–100.

Gao, H., M. Lemmon, and K. Li (2012), Is ceo pay in us public firms efficient? new evidence from private firms, working Paper.

Geletkanycz, M. A., B. K. Boyd, and S. Finkelstein (2001), The strategic value of ceo external directorate networks: Implications for ceo compensation, *Strategic Management Journal*, *22*(9), 889–898.

Goldberg, P. K., and N. Pavcnik (2007), Distributional effects of globalization in developing countries, *Journal of Economic Literature*, *45*(1), 39–82.

Gourio, F., and A. Kashyap (2007), Investment Spikes: New Facts and a General Equilibrium Exploration, *Journal of Monetary Economics*, *54*, 1–22.

Gulisashvili, A. (2012), *Analytically Tractable Stochastic Stock Price Models*, Springer.

Hartzell, J. C., and L. T. Starks (2003), Institutional investors and executive compensation, *Journal of Finance*, *58*(6), 2351–2374.

Helpman, E., M. J. Melitz, and S. R. Yeaple (2004), Export versus fdi with heterogeneous firms, *American Economic Review*, *94*(1), 300–316.

Helpman, E., O. Itskhoki, and S. Redding (2010), Inequality and unemployment in a global economy, *Econometrica*, *78*(4), 1239–1283.

Hennessy, C. A., and T. M. Whited (2005), Debt dynamics, *The Journal of Finance*, *60*(3), 1129–1165.

Heston, A., R. Summers, and B. Aten (2002), Penn world table, *Center for International Comparisons at the University of Pennsylvania*.

House, C. (2008), Fixed Costs and Long-Lived Investments, *NBER Working Paper, No.14402*.

Hummels, D. (2007), Transportation costs and international trade in the second era of globalization, *Journal of Economic Perspectives*, *21*(3), 131–154.

Internal Revenue Service (2013a), Publication 15, employer's tax guide, circular E.

Internal Revenue Service (2013b), Publication 15-a, employer's supplemental tax guide, circular E.

Internal Revenue Service (2013c), Publication 15-b, employer's tax guide to fringe benefits, circular E.

Jacks, D. S., C. M. Meissner, and D. Novy (2008), Trade costs, 1870-2000, *American Economic Review*, *98*(2), 529–34.

Jensen, M. (1997), Eclipse of the public corporation, *Harvard Business Review (Sept.-Oct. 1989), revised*.

Khan, A., and J. Thomas (2003), Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?, *Journal of Monetary Economics*, *50*(2), 331–360.

Khan, A., and J. Thomas (2007), Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies, *The American Economic Review*, *97*(4), 1165–1188.

Khan, A., and J. Thomas (2008), Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics, *Econometrica*, *76*(2), 395–436.

Krusell, P., and A. Smith (1997), Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns, *Macroeconomic Dynamics*, *1*(02), 387–422.

Krusell, P., and A. Smith (1998), Income and Wealth Heterogeneity in the Macroeconomy, *The Journal of Political Economy*, *106*(5), 867–896.

Lucas, R. E., Jr (1978), On the size distribution of business firms, *Bell Journal of Economics*, *9*(2), 508–523.

Manasse, P., and A. Turrini (2001), Trade, wages, and 'superstars', *Journal of International Economics*, *54*(1), 97–117.

McCallum, A. H. (2013), The structure of export entry costs, university of Michigan Econonomics Department mimeo.

Meckl, J., and B. Weigert (2011), Occupational choice, aggregate productivity, and trade, *IZA Discussion Papers 5497*, Institute for the Study of Labor (IZA).

Melitz, M. (2003), The impact of trade on intra-industry reallocations and aggregate industry productivity, *Econometrica*, *71*(6), 1695–1725.

Miao, J., and P. Wang (2011), A Q-Theory Model with Lumpy Investment, *Mimeo*.

Monte, F. (2011), Skill bias, trade, and wage dispersion, *Journal of International Economics*, *83*(2), 202–218.

Murphy, K. J. (1999), Executive compensation, in *Handbook of Labor Economics*, vol. 3, Part B, edited by O. C. Ashenfelter and D. Card, pp. 2485 – 2563, Elsevier, doi:http://dx.doi.org/10.1016/S1573-4463(99)30024-9.

Piketty, T., and E. Saez (2003), Income inequality in the united states, 1913-1998, *The Quarterly Journal of Economics*, *118*(1), 1–39.

Poterba, J. M., and D. R. Feenberg (2000), The income and tax share of very high-income households, 1960-1995, *American Economic Review*, *90*(2), 264–270.

Resnick, S. I. (1987), *Extreme Values, Regular Variation, and Point Processes*, Springer.

Richardson, H. (1995), Control your costs then cut them, *Transportation & Distribution*, *36*(12), 94–94.

Roberts, D. R. (1956), A general theory of executive compensation based on statistically tested propositions, *The Quarterly Journal of Economics*, *70*(2), 270–294.

Rose, N., and C. Wolfram (2002), Regulating executive pay: Using the tax code to influence chief executive officer compensation, *Journal of Labor Economics*, *20*(2), S138–S175.

Rotemberg, J. J., and M. Woodford (1991), Markups and the business cycle, in *NBER Macroeconomics Annual 1991, Volume 6*, NBER Chapters, pp. 63–140, National Bureau of Economic Research, Inc.

Slemrod, J. (1996), High-income families and the tax changes of the 1980s: The anatomy of behavioral response, in *Empirical Foundations of Household Taxation*, NBER Chapters, pp. 169–192, National Bureau of Economic Research, Inc.

Tauchen, G. (1986), Finite state markov-chain approximations to univariate and vector autoregressions, *Economics Letters*, *20*(2), 177–181.

Thomas, J. (2002), Is Lumpy Investment Relevant for the Business Cycle?, *The Journal of Political Economy*, *110*(3), 508–534.

Vavra, J., and D. Berger (2012), Consumption Dynamics During the Great Recession, *2012 Meeting Papers 109*, Society for Economic Dynamics.

Veracierto, M. L. (2002), Plant-level irreversible investment and equilibrium business cycles, *The American Economic Review*, *92*(1), pp. 181–197.

Yakovenko, V. M., and A. C. Silva (2005), Two-class structure of income distribution in the usa: Exponential bulk and power-law tail, in *Econophysics of Wealth Distributions*, pp. 15–23, Springer.

Yeaple, S. R. (2005), A simple model of firm heterogeneity, international trade, and wages, *Journal of International Economics*, *65*(1), 1–20.