Nominal Devaluations and Inequality*

Andrés Blanco † Andrés Drenik ‡ Emilio Zaratiegui§

September 11, 2020

Abstract

We study the distribution of labor income during large devaluations. Across countries, mean real labor income and inequality falls after large devaluations. To understand inequality dynamics, we use a novel administrative dataset to analyze in depth the 2002 Argentinean devaluation. Following individual workers over time, we show that, after an homogeneous fall in labor income, low-income workers experience a faster recovery than high-income workers. Between-firm labor income differences are the main contributors to the heterogeneous recovery. We provide evidence about the role of labor mobility and income floors set by unions for the heterogeneous recovery.

JEL: F31, F41, F44

Keywords: large devaluations, labor income risk, inequality.

*We thank for helpful discussions Ariel Burstein, Javier Cravino, Fatih Guvenen, Loukas Karabarbounis, Chris Moser, Virgiliu Midrigan, Pablo Ottonello, Diego Perez, Felipe Saffie, and Linda Tesar, and seminar participants at University of Michigan, SED, UC Davis, UC Santa Cruz, International Macro-Finance Conference at Chicago Booth, Columbia University, Cleveland FED, Universidad de San Andres, Universidad Torcuato Di Tella, NYU, and Chicago Fed. We thank Javier Tasso for superb research assistance. We are particularly grateful to the Ministry of Production and Labor in Argentina for granting access to the data. All analysis of confidential data was carried out on-site and in accordance with Argentinian confidentiality laws.

†Corresponding author: University of Michigan, jablanco@umich.edu, 611 Tappan Av., Ann Arbor, Michigan, 48109.

‡Columbia University, ad3376@columbia.edu.

§Columbia University, emiliozaratiegui@gmail.com.

1 Introduction

Sudden and large nominal exchange rate (NER) devaluations are associated with significant and abrupt increases in inflation, alongside collapses in output. Despite the importance of these episodes in emerging economies, there is little evidence of their heterogeneous effects on workers. The lack of empirical facts to guide economic research motivates our question: How does the labor income *distribution* evolve during large NER devaluations?

We establish two empirical regularities in the labor market during large NER devaluations: Mean real labor income and inequality fall during these episodes. To establish these facts, we assemble a panel dataset across emerging economies with data on the NER, inflation, output, mean labor income, and inequality. These data allow us to study the labor market dynamics within the broad macroeconomic context across countries. We find that large devaluations are associated with a significant drop and recovery of output, together with an increase in inflation of around one-third of the devaluation rate. Since mean nominal labor income is constant and inflation increases during the year of the devaluation, real labor income falls by 25%. During the recovery, labor income inequality—as measured by the Gini coefficient—drops by 4 points (to illustrate the significance of the decline, the Gini coefficient has increased by 6.4 points over the last 40 years in the U.S.). Finally, we show that recessions without devaluations are associated with stable inflation and real labor income, and increasing inequality. All these facts are not driven by specific episodes, such as devaluations contemporaneous with sovereign defaults or banking crises.

While the cross-country evidence allows us to establish a surprising fact during large devaluations—i.e., the drop in inequality—it does not allow us to understand the reason why inequality falls. To understand the economic mechanisms behind our main fact, we use a novel monthly administrative employer-employee matched dataset that covers the 2002 devaluation in Argentina. We leverage three characteristics of these data: frequency, quality, and coverage. First, we can differentiate between income fluctuations that result from variations in earnings and employment status and their interaction for labor income, since we observe workers and employers at a monthly frequency. The higher frequency, in turn, allow us to precisely capture patterns of labor mobility. Second, the source of the data is employers' sworn statements used for tax purposes and to determine workers' social security contributions. Hence, our data contain little measurement error and no top-coding, which are common problems with survey-based micro-data. Third, our dataset includes the universe of formal workers and firms.

We first document that the dynamics of output, inflation, mean and the Gini coefficient of real labor income during the 2002 devaluation in Argentina follow similar patterns than in the cross-country analysis. The pass-though of the NER to inflation was 28% (30% across countries), the drop of labor income is 26% (25% across countries), and the Gini coefficient declined by 8 percentage points (4 percentage points across countries). Our data allow us to provide a more detailed picture of the dynamics of inequality. During the 2002 devaluation, we find: (i) almost no movement of the income distribution during the two years before the devaluation, (ii) an homogeneous drop in income during the first two quarters after the devaluation, and (iii) a heterogeneous recovery. While the 10th and 25th percentiles of the income distribution recover to their pre-devaluation levels 21 months after the devaluation, it takes 61 months for the 90th percentile to recover.

While the dynamics of different percentiles of the distribution are informative of cross-sectional statistics, they do not necessarily reflect individual income dynamics of workers across the income distribution. We extend the analysis by ranking workers according to their pre-devaluation (2000-2001) income and analyzing their within-worker average income growth. We find an empirical pattern best described as a "parallel drop and pivoting." In the year after the devaluation, there is a parallel average within-worker drop in income of 24% across the pre-devaluation distribution, followed by a clockwise pivoting of the cumulative mean income growth centered around the income growth of the highest-income workers. That is, after four years, workers at the 10th percentile of the pre-devaluation distribution had experienced an average cumulative income growth of 43% relative to the month preceding the devaluation, while the average cumulative growth of those in the 90th percentile was -6%. Thus, there was a clockwise pivoting over time of income growth across the income distribution. Such heterogeneous post-devaluation income growth is linearly decreasing in workers' initial income ranks.[1] Thus, low income workers can better hedge against the increase in inflation.

Between-firm heterogeneity is the main contributor to the "pivoting" effect in the recovery. To reach this conclusion, we decompose the recovery of income across the pre-devaluation income distribution into between-sector, between-firm, and within-firm components. For income levels below (resp. above) the 60th percentile, the average sectoral and workers' income growth—relative to the average of the firm—are almost constant (resp. decline). The recovery of firms' average labor income relative to the sector—the between-firm component—exhibits the largest heterogeneity: The average growth of the between-firm component for workers in the 10th (resp. 90th) percentile was 20% (resp. -8%), and it was linearly decreasing across the percentiles of the distribution. Thus, our data suggest that to study the heterogeneous labor income dynamics during large devaluations, economists should focus their attention on explaining the drivers of firms' average labor income relative to the sector. In conclusion, firms play a critical role for the decline in inequality during the recovery of labor income.

Having established the main empirical facts, we next provide evidence of driving mechanisms. Given the importance of between-firm heterogeneity for the "pivoting" effect in

[1] This fact is robust to further splitting workers according to their characteristics before the devaluation (e.g., age group, 1-digit industry, gender, full-time status, pre-devaluation trends) and to the inclusion of workers with zero monthly income in the formal private sector.

the recovery, we study the contribution of labor mobility across firms.[2] Similarly, since the between-sector and between-worker components play a role in the slower recovery of high-income workers, we study the importance of income floors set by unions across sectors and occupations.

We find that the primary adjustment channel driving the "parallel drop and pivoting" pattern in the data is labor mobility. We demonstrate the claim in three steps. First, we show that the cumulative probability of separations and job-to-job transitions are decreasing in income. Thus, labor mobility is more prevalent among low-income earners. Second, we show that average income growth across jobs after a separation is only positive for low-income earners, and it is positive and decreasing in pre-devaluation income after a job-to-job transition. In the last step, we perform an accounting exercise by constructing several counterfactual income series to evaluate the quantitative role of labor mobility. We construct cumulative labor income changes without considering income changes experienced after separations, job-to-job transitions, or both. We show that workers in the 10th (resp. 90th) percentile of the income distribution experienced a 10% (resp.-5%) faster recovery in the data relative to the counterfactual series that exclude changes in income after separations and job-to-job transitions. The quantitative magnitude of the recovery in the counterfactual labor income series is one-half of the recovery in the between-firm component of labor income. Thus, labor mobility is a main economic mechanism that allowed low-income workers to hedge against inflation.

The dynamics of income floors set by unions is another important mechanism. To demonstrate this, we perform two analyses. First, we digitize the wage scales in collective bargaining agreements (CBA from hereon) in sectors with strong unions and broad coverage (those sectors employ 18% of workers in the sample) to study the income dynamics by unionization status. We find that the income growth of unionized workers with incomes close to the CBA-mandated floors is 30% higher than non-unionized workers. In those sectors, unions negotiated an increase in income between 30% and 60% above inflation. Second, across all sectors, unionized workers are mainly middle-income earners, and their income recovers 6% more than non-unionized workers.

We also find that international trade and income risk have a limited role in explaining the central fact. Given the large change in relative prices induced by the devaluation, we study labor income in tradable and nontradable sectors separately. We find that the NER and sectoral labor income are correlated, and their correlation is a function of trade exposure. Relative tradable income was decreasing before the 2002 devaluation, but it reverted after the devaluation–relative tradable income persistently increased by 10% relative to the nontradable sector. Despite these findings, trade exposure cannot explain the decline in

[2] The higher prevalence of labor mobility among low-income workers has been documented in the US by Karahan, Ozkan and Song (2019). Here, we show that similar mobility patterns have implications for the distributional impacts of large devaluations.

inequality, because it is not the case that low (resp. high) income households are mostly employed in tradable (resp. nontradable) sectors that benefited from the devaluation. Finally, we explore the possibility that lower inequality is the result of less volatile income growth after the devaluation. Intuitively, the dispersion of income after the devaluation is a function of the dispersion of changes in income. Changes in income risk cannot be a driver of the decline in inequality, as the dispersion of year-over-year income growth increased significantly following the 2002 devaluation (i.e., the interquartile range increased by 20%).

Finally, we study the role of policy changes and additional dimensions of the labor market (such as the informal sector) to better understand the heterogeneous dynamics during devaluations. Since our dataset does not include information on hours of work, we make use of household surveys and data on workers with full-time labor contracts. The decrease in inequality is also observed for full-time workers and in the distribution of hourly wages. Since our analysis is based on real labor income constructed with the aggregate CPI, we also reproduce our central fact with real income constructed using income-specific CPIs. The pivoting of the real income recovery decreases slightly, since pivoting in the income-specific CPIs is quite small. Finally, we analyze the role played by a policy intervention after the 2002 devaluation: minimum wage adjustments. We find that this policy cannot explain the decrease in inequality after the 2002 devaluation.

Literature review. We highlight our contributions to two areas of the literature: (i) the macroeconomic consequences of large devaluations and (ii) real labor income dynamics after a significant increase in inflation.

Our paper advances previous work on the economic consequences of large devaluations. Burstein, Eichenbaum and Rebelo (2005) find that on average, 38% of total nominal exchange rate depreciation is incorporated into CPI prices within 24 months. Thus, large devaluations tend to be followed by large spikes in aggregate inflation. Eichengreen and Sachs (1985) and Schmitt-Grohé and Uribe (2016) offer a more aggregate perspective on the matter through their work on the interaction between the labor market and devaluations. They argue that a devaluation, and its upward pressure on prices, can overcome downward nominal wage rigidities and stimulate output.[3]

Previous literature has measured the distributional effects of monetary and exchange rate policy that originate from different channels. Doepke and Schneider (2006) and Verner and Gyongyosi (2018) study the distributional impact resulting from the revaluation of nominal

[3] Gopinath and Neiman (2014) and Blaum (2019) study the effect of large devaluations in aggregate productivity through fluctuations in input trade. Gopinath and Neiman (2014) find a drop in aggregate imported inputs in the 2002 Argentinian devaluation, and argue for a drop in productivity within the context of their model. Blaum (2019) shows an increase in the imported input share across large devaluations. In a model consistent with the fact that exporters have a larger share of imported inputs, Blaum (2019) shows that aggregate productivity can increase after a devaluation. See also Mendoza (2010), Ates and Saffie (2016), and Benguria, Matsumoto and Saffie (2020) for the study of sudden stops in emerging economies.

debt. Previous research that focuses on the distributional impact of large devaluations has found that low-income workers are more negatively affected. They experience larger increases in household-specific inflation (e.g., Cravino and Levchenko (2017)) and a larger negative revaluation of their nominal assets since they tend to save in local currency assets (e.g., Drenik, Pereira and Perez (2018)). On the other hand, Hausman, Rhode and Wieland (2019) finds that the dollar devaluation contributed to the recovery after the Great Recession through the redistribution to the farming sector. In this paper, we extend the previous analysis by documenting inequality dynamics in a cross-section of large devaluations and focusing on the dynamics of the labor income *distribution* after large devaluations using administrative micro-data from an emerging economy.

Previous work has shown a large asymmetry in the wage change distribution in low- and stable-inflation environments. This fact is interpreted as evidence of downward nominal wage rigidities. These facts are reported in the U.S. and Europe by Kahn (1997), Dickens *et al.* (2007), Sigurdsson and Sigurdardottir (2011), Le Bihan, Montornès and Heckel (2012), Barattieri, Basu and Gottschalk (2014), and more recently by Grigsby, Hurst and Yildirmaz (2019). This paper documents the evolution of Argentina's real income distribution after an increase in inflation of 35%. We provide novel evidence detailing the different speeds at which real labor income adjusts for workers across the income distribution after a significant increase in inflation. We find that high-income earners take four more years to revert to their pre-shock level than low-income earners. That is, the recovery of real income is heterogeneous, can be predicted by workers' characteristics, and has large effects on inequality.

Layout. The paper is organized as follows. Section 2 describes the data. Section 3 presents the aggregate facts in the cross-country analysis of large devaluations. Section 4 revisits those aggregate facts in our main episode of analysis. Section 5 presents evidence on the mechanisms behind these facts. Sector 6 demonstrates the robustness of our findings and Section 7 concludes.

2 Data

This section describes the international data and the novel dataset we leverage to study the dynamics of the income distribution after large devaluations. Interested readers should refer to Online Appendix Section A.1 for a detailed description of the data construction for the cross-country analysis and Online Appendix Section A.3 for a discussion of variables in SIPA, sample construction, and cross-validation of SIPA results.

Data for cross-country analysis. We analyze five variables across 44 countries: output, NER, inflation, real labor income, and a measure of inequality. For output, we use GDP at constant prices in local currency from the World bank. We use the consumer price index to measure inflation and the nominal exchange rate is the exchange rate between the local currency and the U.S. dollar. Inflation and NER data come from the IMF International Financial Statistics Dataset. Real labor income is constructed as the average monthly wage in local currency deflated by the CPI (see Table A.1 for data sources in each country). Finally, we measure inequality with the Gini coefficient provided by the Word Bank. The Gini coefficient is based on household survey data from national statistical agencies and World Bank country departments. The Gini coefficient is mainly computed with data on disposable labor income (see the end of Section 3 for a broader discussion about the Gini coefficient).

Labor income data for the 2002 Argentinean devaluation. We use administrative employer-employee matched monthly panel data from Argentina. The data start in July 1994 and end in June 2019. Our data source is Argentina's national social security system ("Sistema Integrado Previsional Argentino", SIPA from hereon). By law, all employers in the formal sector must present sworn statements providing relevant worker compensation information to SIPA every month.

SIPA tracks each worker's total monthly labor income in the formal sector without measurement error or top-coding, including all forms of payment that could trigger tax liabilities or social security contributions (e.g., base wage, bonuses, overtime compensation, etc.). The dataset also includes relevant demographic information on each worker and their job, as well as some characteristics of the firm, such as 4-digit industry and state. Importantly, SIPA also provides firm and worker identifiers that are consistent across the entire period, which allow us to analyze income dynamics for individual workers and firms at a monthly frequency for up to 26 years.

The dataset covers the universe of formal workers employed in all regions, private industries, types of contracts (internships, temporary workers, full-time employees, etc.), and in the public sector. One of the benefits of analyzing the Argentinian labor market is that relative to other Latin American economies, the informality rate is not as high—e.g., Gasparini and Tornarolli (2009) report a formality rate in Mexico of 45%. Figure D.15-Panel B shows the time series of the share of formal employment in the private sector for male salaried workers aged between 25 and 65 in Argentina. Throughout the sample period the average formality rate was 70%. We conclude that our data cover a large share of the overall population.

When we analyze labor income dynamics in a large devaluation in Argentina, we present facts about the (log) real pretax total labor compensation of male workers aged between 25 and 65 in the private sector.[4] We restrict our sample to male workers aged between 25 and

[4] Due to the intervention of inflation statistics in Argentina in 2007, we use consumer price indices provided by national statistics before 2007 and PriceStats from 2007 onward to construct real labor income. In our

65 years to avoid issues related to labor force participation and retirement. Finally, we drop observations coming from job spells that involve workers employed in the public sector, since their wages might not be market-determined and subject to other nonmarket forces.

We apply some filters to monthly real labor income in our analysis. We eliminate outliers and winsorize top observations. We define outliers as workers who earn less than half of the monthly minimum wage. Because the minimum wage in Argentina has changed over time, we use the 1996 value in real terms (i.e., $200 per month) and adjust it by the average growth rate of real wages in the entire sample (i.e., 2% annual growth). We winsorize observations above the 99.999th percentile. We also omit the first and last wage of each job spell due to time aggregation concerns, since we do not know the day a spell starts/ends or whether the last wage includes severance pay.[5] Although we do not consider these monthly salaries in the analysis of labor income, we use these observations to analyze employment flows. The final dataset with our sample selection and filters contains more than 700 million worker-month observations. Finally, we seasonally adjust all time series using the X-13ARIMA-SEATS seasonal adjustment program developed and used by the U.S. Census Bureau. Since this is one of the few papers that use the SIPA dataset, we provide a further discussion of the quality of the data in Online Appendix A.4 and A.5.

Additional data for the 2002 Argentinean devaluation. We complement the SIPA database with the information contained in Collective Bargaining Agreements (CBAs, from hereon) negotiated by trade unions at the sectoral level. In Argentina, a single union has monopoly power to represent workers at the sectoral level. That union signs a contract covering all workers employed in a specific subset of occupations in the sector, regardless of their membership status. We digitize these contracts at the sectoral level for several of the most important unions (those sectors employ 18% of workers in the sample). We also use data from the Permanent Household Survey ("Encuesta Permanente de Hogares", EPH from hereon), which is the main household survey in Argentina.

3 Two Facts After Large NER Devaluations

What are the empirical regularities about the labor income distribution during large NER devaluations? We find that during these episodes mean labor income drops by 25%, and the Gini coefficient falls by 4 points four years after a devaluation. Since large NER devaluations are associated with meaningful recessions, we revisit these facts during recessions without

baseline analysis, we deflate nominal income with the aggregate CPI. In Section 6, we verify the robustness of our analysis by computing income-specific levels of prices as in Cravino and Levchenko (2017).

[5] We purge the monthly labor income of the 13th salary paid in June and December to avoid spurious seasonality. This extra salary is mandated by law and equals one half of the highest wage paid over the semester. Because we only observe total income before 2008, we use the formula established by law to calculate each worker's 13th salary.

devaluations. We find that during recessions without devaluations mean labor income is constant and the Gini coefficient increases by 2 points. These facts are not driven by specific episodes or special types of devaluations or recessions, such as sovereign defaults or banking crises.

We follow the definition of currency crises by Laeven and Valencia (2012) to identify large NER devaluations. Laeven and Valencia (2012) define currency crises as a nominal depreciation rate of the currency vis-à-vis the U.S. dollar of at least 30%, that is also at least 10% higher than the depreciation rate in the previous year. This definition follows the pioneering work of Frankel and Rose (1996). The sample of large NER devaluations with complete data on the Gini coefficient and labor income includes 19 episodes. We classify a recession without a devaluation to two consecutive drops of GDP of at least 2% without a large NER devaluation. The sample of recessions without devaluations contains 40 episodes. See Table A.3 for a list of all episodes in the analysis.

Figure 1 – Macroeconomic Facts After Large NER Devaluations

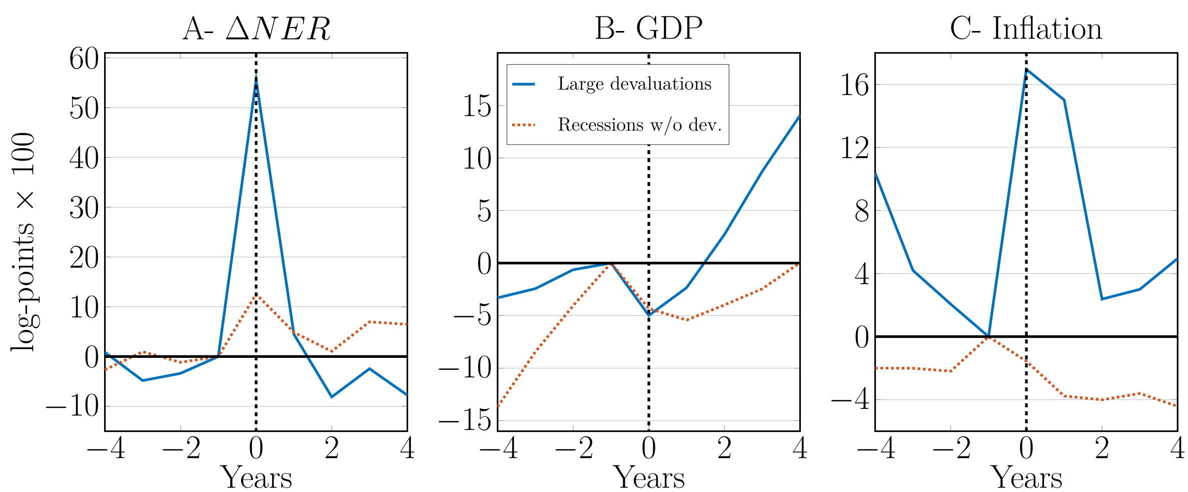

Notes: Panels A to C plot (in the following order) the change in the NER, real GDP, and inflation at an annual frequency. All variables are expressed in log-point × 100 and normalized to zero in year -1. The blue solid line shows the average macroeconomic dynamics in a 8-year window around a large devaluation. The year zero corresponds to the year of the devaluation. The red dotted line plots the same variables for recessions without devaluations. The year zero corresponds to the year with the first drop in GDP. See main text for the description of the episodes.

Macroeconomic context during large NER devaluations. Figure 1 plots the evolution of the average annual NER devaluation rate, real GDP, and inflation in an 8-year window around large devaluations and recessions without devaluations. Large nominal devaluations are associated with a significant recession, recovery of output, and an increase in inflation. The average GDP drop across episodes is 5%, which coincides with the average output drop during recessions without devaluations. While the drop in GDP is similar across these two types of episodes, the recovery is faster during large devaluations. In addition, during large devaluations, there is a large pass-through into domestic inflation. Burstein *et al.* (2005) documents an average elasticity of annual inflation to a large nominal devaluation of one-third across emerging economies. This number coincides with the pass-through in our sample: The average ratio of annual changes in inflation over annual changes in the NER is 31% (i.e., 17/55 =0.31). On the other hand, during recessions without devaluations, inflation drops relative to its pre-recession level.

Figure 2 – Labor Market Facts After Large NER Devaluations

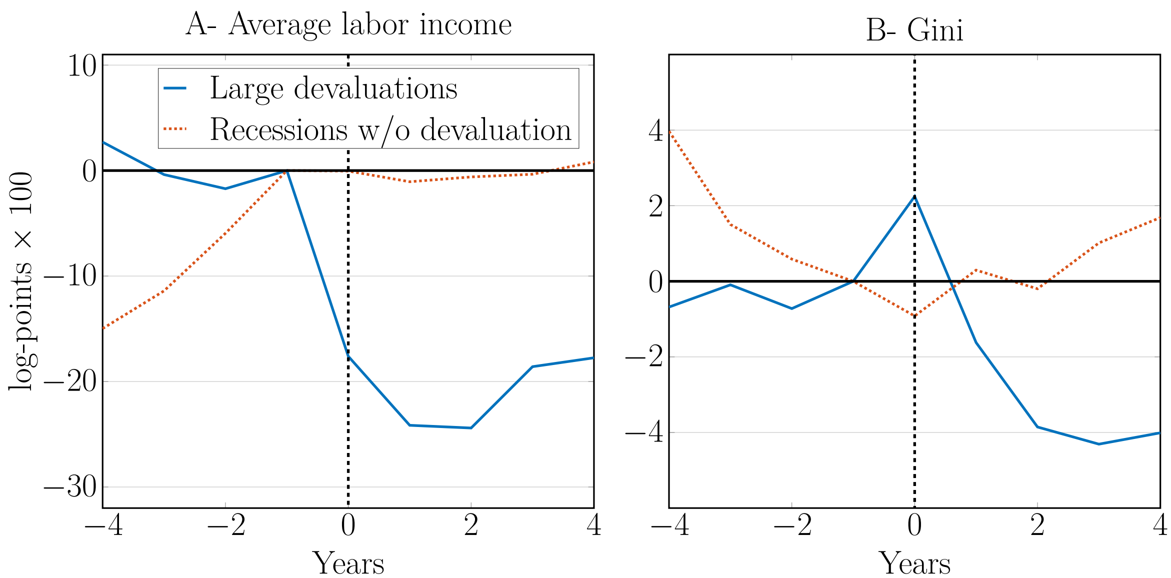

Notes: Panels A and B plot (in the following order) the average labor income and Gini coefficient. Average labor income is expresses in log-points × 100 and normalized to zero in year -1. The blue solid line shows the average dynamics in a 8-year window around a large devaluation. The red dotted line plots the same variables for recessions without devaluations. The year zero corresponds to the year with the first drop in GDP. The episodes included in large devaluations and recessions without devaluations are the same as in Figure 1.

Labor market facts during large NER devaluations. During large devaluations, real labor income falls by 25% and the Gini coefficient falls by 4 points below its pre-devaluation level. Figure 2 plots the average labor income and the Gini coefficient following the same format as Figure 1. The figure shows no pre-trends in mean labor income and the Gini coefficient before devaluation episodes. During the devaluation, *nominal* labor income is constant, thus *real* labor income falls by the same magnitude than the increase in inflation. One year after a large devaluation, real labor income drops by less than the increase in inflation and then starts recovering two years after the devaluation. While mean real labor income falls during large devaluations, we do not find this pattern in recessions without devaluations since nominal and real labor income are almost constant during these episodes.

The Gini coefficient falls when real income recovers. Four years after the devaluation, the Gini coefficient is 4 points lower than its pre-devaluation level. The fall in inequality measured by the Gini is significant. To illustrate this, in a country where income inequality has received considerable attention in academic and political circles such as the U.S., the Gini coefficient has increased by 6.4 points over the last 40 years.

The facts behind Figure 2 are surprising, and the goal of the rest of the paper is to explore the economic mechanisms behind them. They are surprising because they show that the main source of income for a large majority of the population, i.e. labor income, becomes less unequally distributed during nominal devaluations. We do not study the consequences nor the policy implications of lower inequality. Instead, to the best of our knowledge, our contribution is to document for the first time the dynamics of the *distribution* of labor income after devaluations and present evidence of the mechanisms driving them. Before exploring the mechanisms at play during large devaluations with the administrative labor income data from Argentina, we discuss our measure of inequality and the robustness of our facts to other confounding factors.

Measures of income inequality. Given the lack of readily-available quality administrative labor income data across emerging economies, we rely on the World Bank's Gini coefficient to establish an empirical regularity about labor income inequality during devaluations. There are several advantages of this measure. First, its frequency is annual, thus we can study its evolution within a 8-year window after large devaluations. Second, while in principle the Gini coefficient is constructed with consumption or income data, depending on the country, in our sample of large devaluations (resp. recession without devaluation), in 14 out of 19 (in 35 out of 40) episodes inequality is computed with income data and the rest with consumption data. Third, in principle, the World Bank's objective is to measure inequality of total income. In practice, the Gini coefficient mostly captures labor income inequality for two reasons: i) actual lack of capital income for the majority of households in emerging economies, and ii) the focus on labor income (or lack of data) in household surveys. Finally, the Gini coefficient measures household income per capita, i.e., it is assumed that all households members receive the same share of household income. These data has been previously used in the literature (see, e.g., Pinkovskiy and Sala-i Martin, 2016). Similarly, these data are one of the data sources behind the World Income Inequality Database (developed by United Nations and used by, e.g., Young, 2013, Fajgelbaum and Khandelwal, 2016).

Robustness. Given the sample size in our list of episodes, a detailed multivariate analysis that controls for differences across episodes would not be feasible. However, in the spirit of showing that the aggregate facts are not driven by particular devaluations or special kinds of recessions, we reproduce the main graphs for different subsets of episodes. In Online

Appendix A.2, we show that similar patterns are observed when we consider episodes: i) episodes that (do not) coincide with banking crisis, ii) episodes without sovereign defaults, iii) episodes in which inequality measures are based on households' income—and not consumption, iv) episodes without hyperinflation, v) episodes with short recessions, and vi) episodes that occurred from the year 2000 onwards. Although there are quantitative differences across sub-samples, we consistently find that large devaluations are followed by declines in average real labor income and inequality.

4 Revisiting the Facts in Argentina

This section uses the novel microdata on monthly labor income to revisit the previous section's empirical regularities in the 2002 Argentinean devaluation. We find similar qualitative patterns in our data: i) output exhibited a significant drop and recovery, ii) inflation increased by one-third of the change in the NER, iii) mean labor income dropped by the same amount of the increase in inflation in the first year, and iv) the Gini coefficient declined when mean labor income recovered. We finish this section with a deeper discussion of inequality dynamics based on cross-sectional moments. Across all the different inequality measures in the cross-section, inequality falls during the recovery of real income after the devaluation mainly because the bottom of the income distribution recovers faster than the top.

Macroeconomic context. The 2002 Argentinean devaluation presents similar dynamics than the cross-country analysis for output, inflation, labor income, and inequality. Figure 3-Panel A shows year-over-year inflation and nominal exchange rate growth and Figure 3-Panel B shows the (log) real quarterly output. Figure 3-Panels C and D show average real labor income and the Gini coefficient at a monthly frequency. We mark the recession period in gray and the month of the devaluation with a dotted black vertical line. To contextualize our measurement exercise, we first describe the macroeconomic environment during the period of analysis (i.e., 1997-2007).

Between 1997 and 2007, there were two exchange rate regimes: a fixed exchange rate from January 1994 to December 2001, in which the national currency was pegged one-to-one to the U.S. dollar, and a floating exchange rate from January 2002 to the present. In the first month of 2002, Argentina abandoned its one-to-one peg to the U.S. dollar. The resulting devaluation rate was 120% (in log points). The size of the devaluation took market participants by surprise.[6]

[6]In Appendix B.1, we present data on exchange rate expectations from a survey of professional forecasters provided by Consensus Economics. In December 2001, professional forecasters were expecting a devaluation of 7% within the following 12 months, so clearly, a devaluation rate of more than 100% had a sizable unexpected component. In Appendix B.2, we plot the dynamics of output per worker as a simple measure of labor productivity.

Figure 3 – Labor Market Facts after the 2002 Argentinean Devaluation

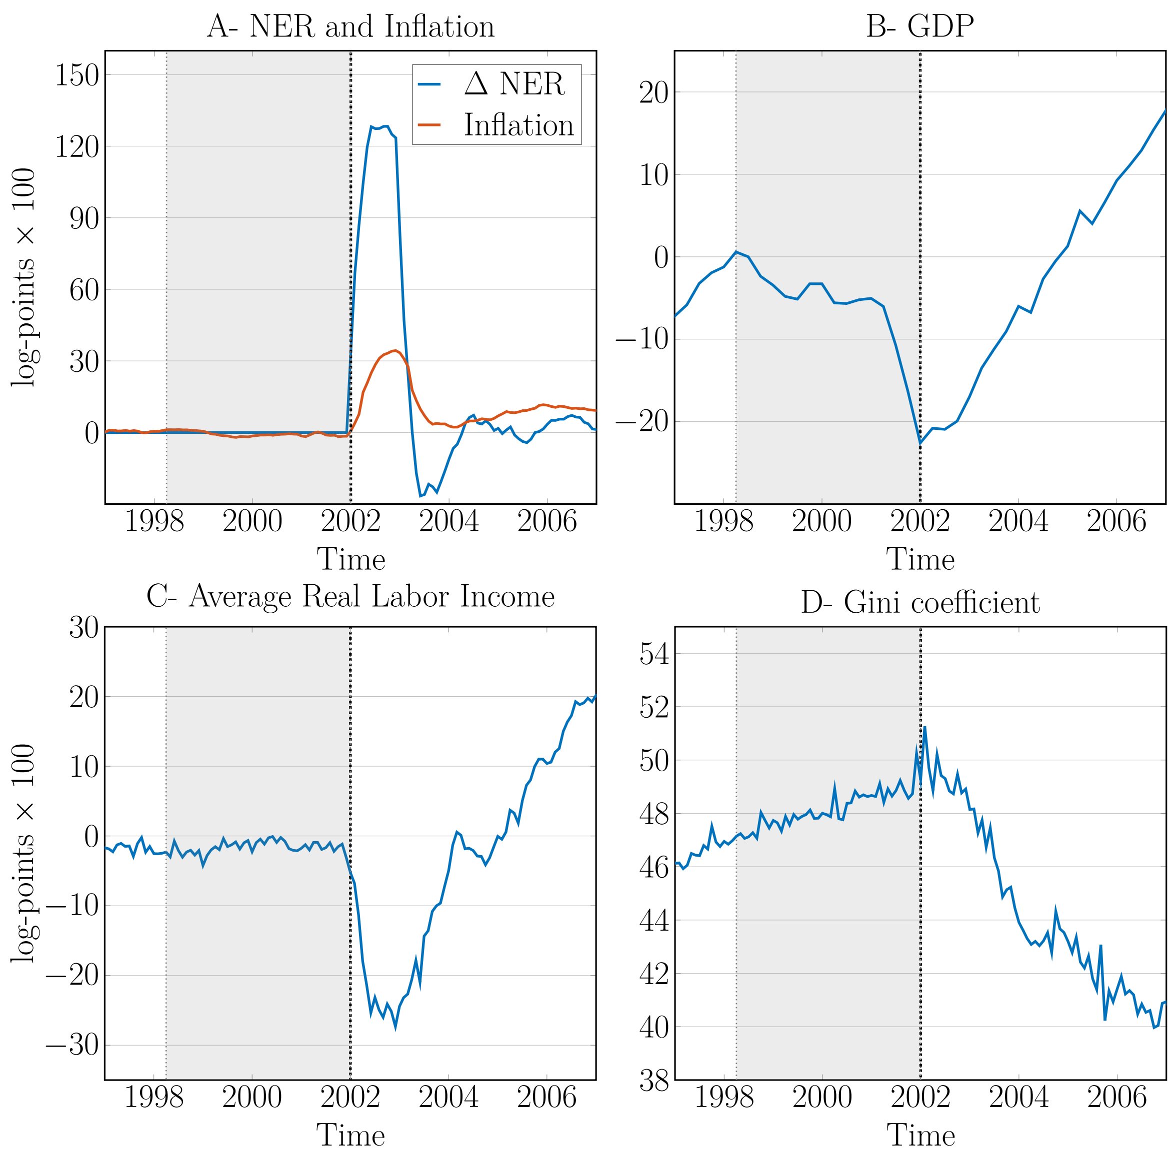

Notes: The figure plots four macroeconomic and labor market time series in Argentina for the period between 1997 and 2007. Panel A plots the NER (blue) and inflation (red), and Panel B plots the real GDP. Panel C shows the average real labor income and Panel D the Gini coefficient. All variables are expressed in log-points × 100. GPD is computed at a quarterly frequency, seasonally adjusted, and normalized to zero in the third quarter of 1998. Inflation, NER, average labor income, and the Gini coefficient are computed at monthly frequency. Real labor income is normalized to zero in the first month of 1996. Recession periods are in gray, and monthly devaluations larger than 30% are marked with dotted black lines.

The 2002 devaluation episode is associated with a significant increase in aggregate prices and the end of the 1998-2002 recession, as in our cross-country analysis. Concerning the price level, the ratio of cumulative logarithmic changes (relative to one month prior to the devaluation) in the price level to cumulative changes in the NER is 0.28 (consistent with the average pass-through measured by Burstein *et al.*, 2005). Concerning the output level, the

1998-2002 recession featured a cumulative output drop of -21% and a limited depreciation of the RER due to a lower inflation rate relative to the US, while the NER remained fixed.

Labor market facts. During the five years before the 2002 devaluation, with a cumulative output drop of 21%, the average labor income remains almost constant. In the first six months after the 2002 devaluation, there was a drop in log average labor income of 26%, as in the international evidence. After this significant drop, it took two years for average income to revert to its pre-devaluation level.

Inequality, measured by the Gini coefficient of the (log) monthly income, increased during the recession, it reached the pick during the devaluation, and started to decline afterward, as we found in the international evidence. The drop in the Gini coefficient accelerated in 2003, when real income started to recover. While the Gini coefficient is useful to establish the empirical fact across countries, we can leverage the microdata in Argentina to analyze other moments of the labor income distribution.

Figure 4 plots moments of the income distribution—normalized percentiles, the interquartile range, and the standard deviation—during five years before and after the devaluation. The first important observation is that, as we can see in the figure, there is no significant fluctuation across percentiles of the income distribution before the 2002 devaluation despite the severity of the recession. This lack of large fluctuations is also reflected in the evolution of the interquartile range and the standard deviation. Second, there is a *homogeneous* drop of 26% across the distribution of real income during the first two quarters after the devaluation. This drop is the result of the rapid increase in inflation and a lack of nominal adjustment of wages.

Despite this homogeneous drop, Figure 4 shows the significant *heterogeneity* in the speed of recovery of real income across different parts of the distribution. While percentiles below the median start recovering after the third quarter, percentiles above the median continue to fall for two additional quarters. Alternatively, note that the 10th percentile of the income distribution recovers to its pre-devaluation level in 21 months, while it takes 61 months for the 90th percentile to recover. This faster recovery of the bottom of the income distribution implies that the distribution became less unequal after the devaluation.

The compression in the distribution during the recovery is reflected in the evolution of the interquartile range and the standard deviation. The interquartile range drops from close to 100% to 80%, and the standard deviation from 79% to 68%. This recovery can be more easily seen in Figure 5, which compares the real income distributions in 2001 and 2006. Four years after the devaluation, there is a substantial shift upward in the bottom of the real income distribution and a compression of real wages from the top.

Figure 4 – Moments of the Distribution of Labor Income

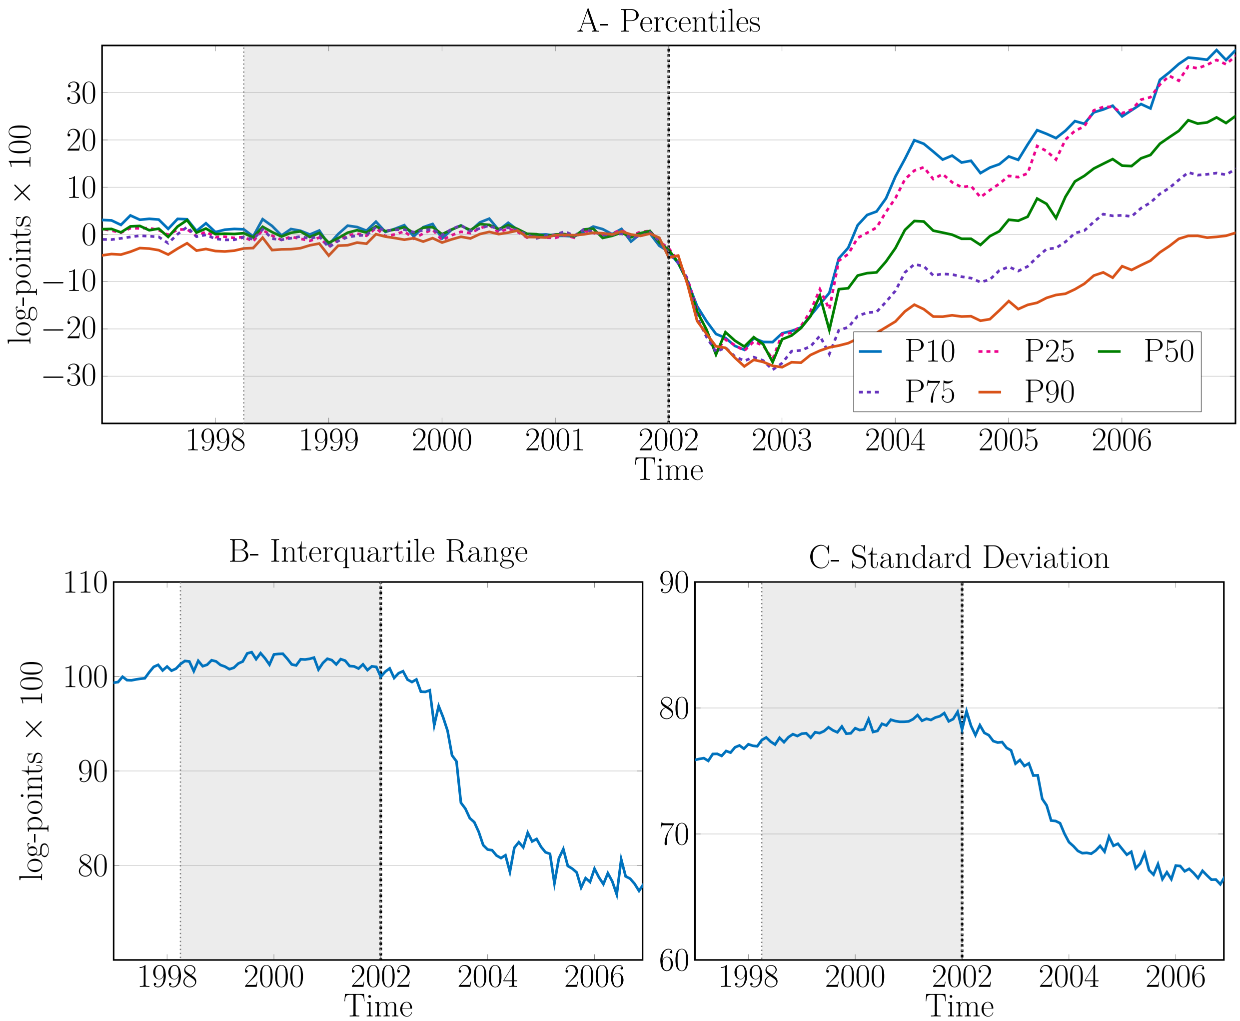

Notes: The figure plots moments of the distribution of monthly real income from January 1997 to December 2007. Panel A plots the percentiles of the log income distribution ($\times$ 100) normalized by their average during 2001. We use Px to denote the x-th percentile of the distribution. Panels B and C plot the interquartile range ($P75 - P25$) and the standard deviation for the same period. Recession periods are in gray, and monthly devaluations larger than 10% are marked with dotted black lines.

5 Why does Labor Income Inequality Fall during Large NER Devaluations?

This section explores the mechanisms behind the fall in inequality during large devaluations. With this goal in mind, we proceed in three steps. In the first step, we study workers' income dynamic conditional on their pre-devaluation income. We find that low-income workers recover from the drop in real income faster than high-income workers—which we label as the "pivoting effect". In the second step, we compute a between-sector, between-firm, and between-workers "variance decomposition" of the conditional income growth. We find that the between-firm component is the main contributor to the pivoting effect. Based on this

Figure 5 – Income Distribution in 2001 and 2006

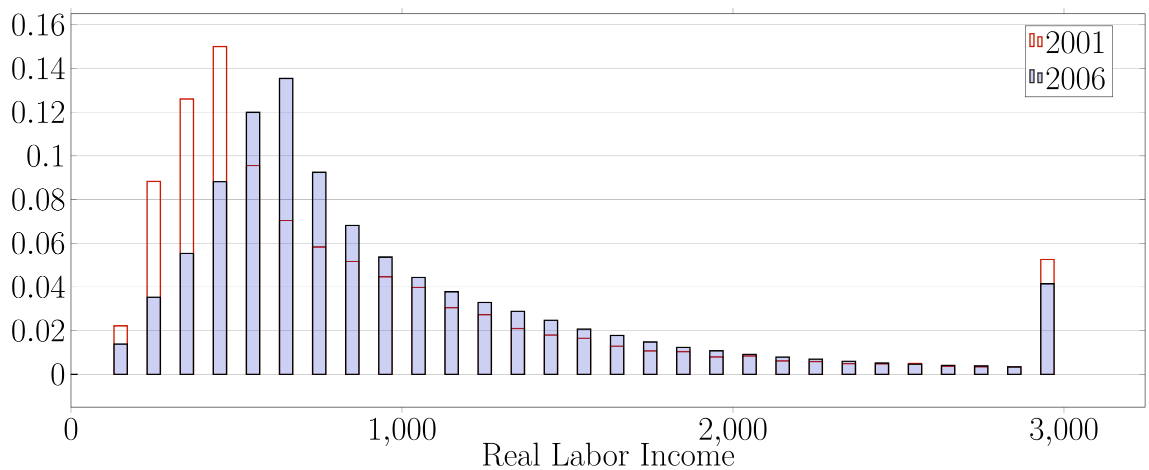

Note: The figure plots the income distribution in 2001 and 2006. Distributions are winsorized using the 95th percentile of distribution as the upper bound.

result, we explore the role of labor mobility across firms in compressing the income distribution. We find that labor mobility can account for half of the heterogeneous recovery of the between-firm component. Given the slower recovery of the between-sector and -worker componets at the top, we explore the role of different income floors set by unions. We find that heterogeneous unionization rates, that lead to different income floor across workers, contributed to the decline in inequality.

5.1 Workers' Income Growth conditional on Income Level

While the analysis in Sections 3 and 4 is informative of cross-sectional statistics, it does not reflect the income dynamics of individual workers across the income distribution during and after the devaluation. This is simply because the identities of the workers within each percentile can change drastically over time. We address this issue by studying workers' income growth conditional on their pre-devaluation level of income.

To do this, we rank workers according to their permanent real monthly income during the pre-devaluation period and group them in percentiles according to this ranking. However, the presence of an age profile in income will render this ranking more favorable toward older workers, thus confounding income and age differences. We address this issue following Guvenen, Ozkan and Song (2014). We first run a pooled regression with all of the data in the sample of log labor income on a set of age and year dummies. Then, we rank workers according to their average log income net of the life-cycle profile during the two years before the devaluation. We drop workers with less than six months of employment during the period 2000-2001, since we cannot precisely capture their average income over the period. Figure 6 shows the mean year-over-year growth of real income (net of the life-cycle profile) from December 2001 onward on the y-axis and the percentiles of the permanent income (net of the life-cycle profile) on the x-axis.[7]

Figure 6 – Average Income Growth Conditional on Average Income in 2000-2001

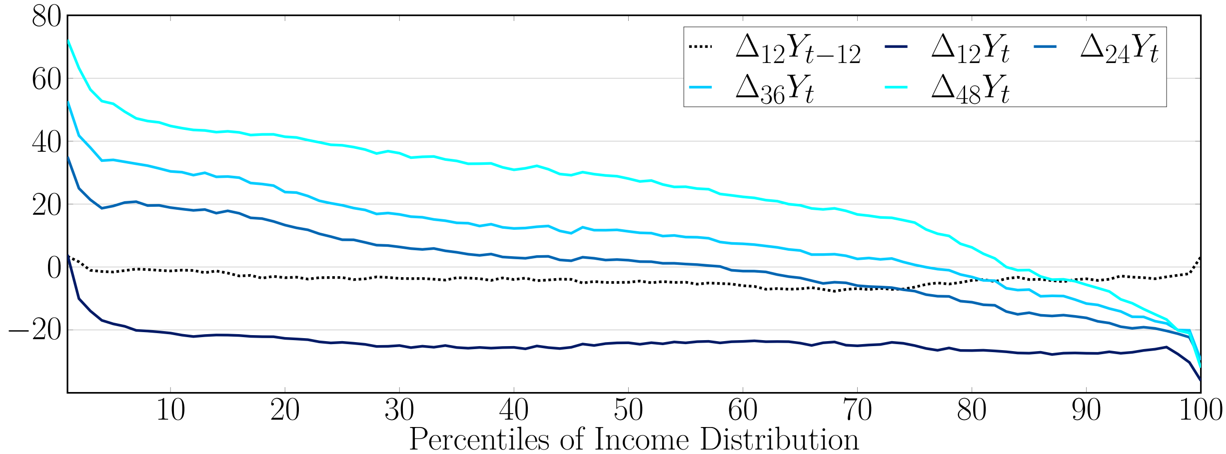

Notes: The figure describes average income growth conditional on the percentile of the distribution of average monthly real income during 2000-2001. The sample is restricted to workers who had at least 6 months of employment during the 2000-2001 period.

The first result is that during the year before the devaluation, the average year-over-year income growth ($\Delta_{12}Y_{t-12}$), is close to zero for all percentiles.[8] This homogeneous average growth disappears after the devaluation, and the pattern that emerges across the income distribution is of a "parallel drop and pivot." That is, in the year after the devaluation, there is a parallel average drop in income ($\Delta_{12}Y_t$) of 24% across percentiles, followed by a pivoting of the cumulative mean income growth centered around the income growth of the highest-income workers. The gap is quantitatively significant. After four years ($\Delta_{48}Y_t$), the average income growth of workers in the 10th percentile of the pre-devaluation distribution had experienced an average cumulative income growth of 43% relative to the month preceding the devaluation, while the average cumulative growth of those in the 90th percentile was -6%.

We extract three conclusions from this analysis. First, income dynamics monotonically depend on the worker's position in the pre-devaluation income distribution. Second, the

[7]Formally, we define the permanent component of income net of the life-cycle profile for agent i as

$$\bar{Y}_t^i \equiv \sum_{m=0}^{23} e^{\tilde{y}_{t-m}^i} \times \mathbb{1}\{N_{t-m}^i = 1\} / \left[\sum_{m=0}^{23} e^{d_{a-m}} \times \mathbb{1}\{N_{t-m}^i = 1\}\right],$$

where t corresponds to the month prior to the devaluation, $\tilde{y}_t^i$ is the log real labor income, d_a are the coefficients of the age dummies in the pooled regression, and N_{t-m}^i is an indicator of employment in period $t-m$. We scale the age dummies so that the fixed effect of a 25 year old worker matches the average labor income of a 25 year old worker in the regression sample.

[8]We use the notation $\Delta_z Y_t = Y_{t+z} - Y_t$.

asymmetric recovery and the decline in inequality are the result of the larger *within-worker* average growth rates for workers at the bottom of the distribution. Third, as we show in Figure B.2 in the Appendix, after the 2002 devaluation, there was a decrease in the labor share from 40% to 31% due to the rapid increase in the inflation rate and the lack of similarly rapid adjustment of nominal labor income—implicitly, a redistribution from the workers to the firms. This sections shows that the redistribution from the firms to the workers, during the recovery of real labor income, is faster at the bottom of the income distribution.

Robustness. To investigate the robustness of the results behind Figure 6, we performed similar analyses using different subsamples of the data. In each case, we found that the main finding on the heterogeneous recovery of real income after the 2002 devaluation still holds. We present our results in Online Appendix Section C.1.

First, we explore the possibility that a subgroup of workers drives the main aggregate result. To address this, we perform additional splits of the data. Given the large change in relative price across sectors brought about by the devaluation, the observed pattern could be the result of a compositional effect. Although we will explore this further below, we reproduce the main finding by splitting the sample according to the 1-digit sector of employment of each worker in December 2001. Figure C.1 shows that the qualitative pattern is present in each of the broad sectors. Similar compositional effects might arise due to differences in the growth rates of income by age. Figure C.2 reproduces the main figure by groups of workers according to their age in December 2001 (25-29, 30-34, etc.) and shows similar patterns in each subgroup of workers. We also reproduce the figure using data on women (see Figure C.3) and find similar results. We also verify that our finding is not determined by the way we construct the measure of permanent income. Thus, following Guvenen *et al.* (2014), we compute the measure of permanent income as the average monthly income for the 5 years prior to the devaluations (as opposed to 2 years, as in the baseline analysis). Figure C.4 shows the results, which are quite similar to those found in the baseline analysis. Finally, we check that the results are not driven by potentially different dynamics of income during the month of December by computing income growth using the average monthly income within the last quarter of the year (see Figure C.7).

One potential concern would be that this fact is driven by changes in the intensive margin (the number of hours worked) or the extensive margin (the employment status of a worker). To address the first concern, we exploit information on the full-time/part-time status of the worker's job. Figure C.5 reproduces the main fact using data on full-time jobs only and shows a similar pattern as in the baseline analysis. To address the second concern, we extend the sample to include the "zeros": If a worker is not employed in the private formal sector in any given month, we replace his income with zero. This generates a balanced panel for each worker employed in December 2001. Figure C.6 shows that the main finding is robust.[9]

Another potential concern of this analysis is that the observed "pivoting" might be the result of mean reversion of labor income. While this concern is qualitatively valid, it is not valid quantitatively for the observed persistence of labor income. We verify this statement by replicating Figure 6 starting in 1997, when aggregate labor income was stagnant, to isolate the effect of mean reversion (see Figure C.9 in the Online Appendix). The patterns between in the two figures are clearly different. In the analysis starting in 1997, average income growth is muted, and there is no "pivoting" effect across the income distribution. In addition, following Guvenen *et al.* (2014), we directly control for different pre-devaluation income growth rates (in addition to controlling for age and the level of income), and find that controlling for past income growth has almost no effect on Figure 6 (see Figure C.10). From this analysis, we conclude that income dynamics after the devaluation are not an artifact of mean reversion and depend on the worker's position in the pre-devaluation income distribution.

5.2 The Role of Sectors, Firms, and Workers for the Pivoting Effect

Can the decrease in inequality be explained by between- or within-group dynamics? This is an important question, since devaluations are associated with large changes in relative prices across sectors and firms, and thus could affect particular groups of workers differently. In Online Appendix Section C.2, we perform a variance decomposition analysis to decompose the overall cross-sectional variance of log real income into between and within components across sectors and firms (see, for example, Song, Price, Guvenen, Bloom and Von Wachter, 2018). There, we find that each of the between-sector, -firm, and -worker components almost equally account for 33% of the decline in labor income inequality.

Although the variance decomposition is a useful starting point used in the literature, it does not provide a characterization of the relevance of the different components (sector, firm, and worker) for the recovery of workers located in different parts of the income distribution. Therefore, we go beyond the standard variance decomposition and, through a series of counterfactual exercises, document how the sectoral and firm components of income differentially affected workers in different parts of the distribution.

In the first exercise, we gauge the relevance of between-sector heterogeneity across the labor income distribution by asking: How would the dynamics of labor income behave if, in each period, workers had earned the average income in the sector? That is, for each worker we compute $\Delta\bar{Y}_{s(it)}$, where $\bar{Y}_{s(it)}$ is the average income in the 4-digit sector s employing worker

[9] To deal with the log and the zeros, we follow the literature (see, for example, Guvenen *et al.*, 2014) and replace $\mathbb{E}_i(\Delta \log y_{it})$ with $\Delta \log \mathbb{E}_i(y_{it})$, where y_{it} is the real income of the i-th worker in period t. By computing the same statistics in our original sample without the zeros we conclude that the differences at the bottom of the distribution between Figure 6 and Figure C.6 are mostly due to Jensen-inequality effects.

i in period t. By construction, this figure also captures the aggregate average increase in labor income. Figure 7-Panel A plots the results by averaging this counterfactual income growth across workers in each percentile of the pre-devaluation distribution (the ranking of workers is the same as the one used in the baseline Figure 6). The two main findings are (i) heterogeneous sectoral labor income growth does not lead to heterogeneous recoveries for workers below the 60th percentile and (ii) part of the decrease in inequality is due to the slower recovery of average sectoral labor income in sectors employing workers at the top of the distribution.

To measure the contribution of the between-firm component across the income distribution, we ask: How would the dynamics of labor income look if, in each period, workers had earned the average income in the firm (net of the average income paid in the sector)? For this we replace the worker's income growth shown before with the worker's growth in $\bar{Y}_{j(it)} - \bar{Y}_{s(it)}$, which is the average income paid in firm j employing worker i in period t net of the average income paid in the sector s of the firm. Figure 7-Panel B shows that this component is responsible for a large fraction of the "pivoting" observed in the baseline Figure 6. Workers below the 60th percentile of the pre-devaluation income distribution experience positive income growth from the between-firm component, while workers above this percentile experience negative income growth. Thus, the decrease in inequality accounted for by the between-firm component is due to monotonically lower average income growth in firms employing higher-income workers.

Finally, the remaining piece of the decomposition is given by changes in $Y_{it} - \bar{Y}_{j(it)}$, which is a worker's i labor income in period t net of the average income paid in the firm employing him. Figure 7-Panel C plots the average growth of this component across the distribution. Most of the heterogeneity in the within-firm and between-worker component comes from faster income growth for workers below the 10th percentile of the pre-devaluation distribution and slower growth for workers at the top of the distribution.

The main takeaway of this section is that in order to explain the labor income dynamics of individual workers at the top of the income distribution, one must focus more on the between-sector and within-firm components. On the other hand, to understand the dynamics for workers at the middle and bottom of the income distribution, one must focus on the between-firm component.[10]

5.3 Economic Mechanism I: Labor Mobility

Given the importance of firms for the pivoting effect, it is natural to ask: what is the role of heterogeneous mobility patterns across the income distribution for the decline in inequality

[10]In Figure C.8, we perform the same analysis on the subsample of workers who, in December 2001, were employed in firms that had on average (during the 2000-2001 period) at least 10 employees. We show that the "pivoting" effect found in the firm component of income is equally important when excluding smaller firms.

Figure 7 – Decomposition of Average Income Growth Conditional on Average Income in 2000-2001

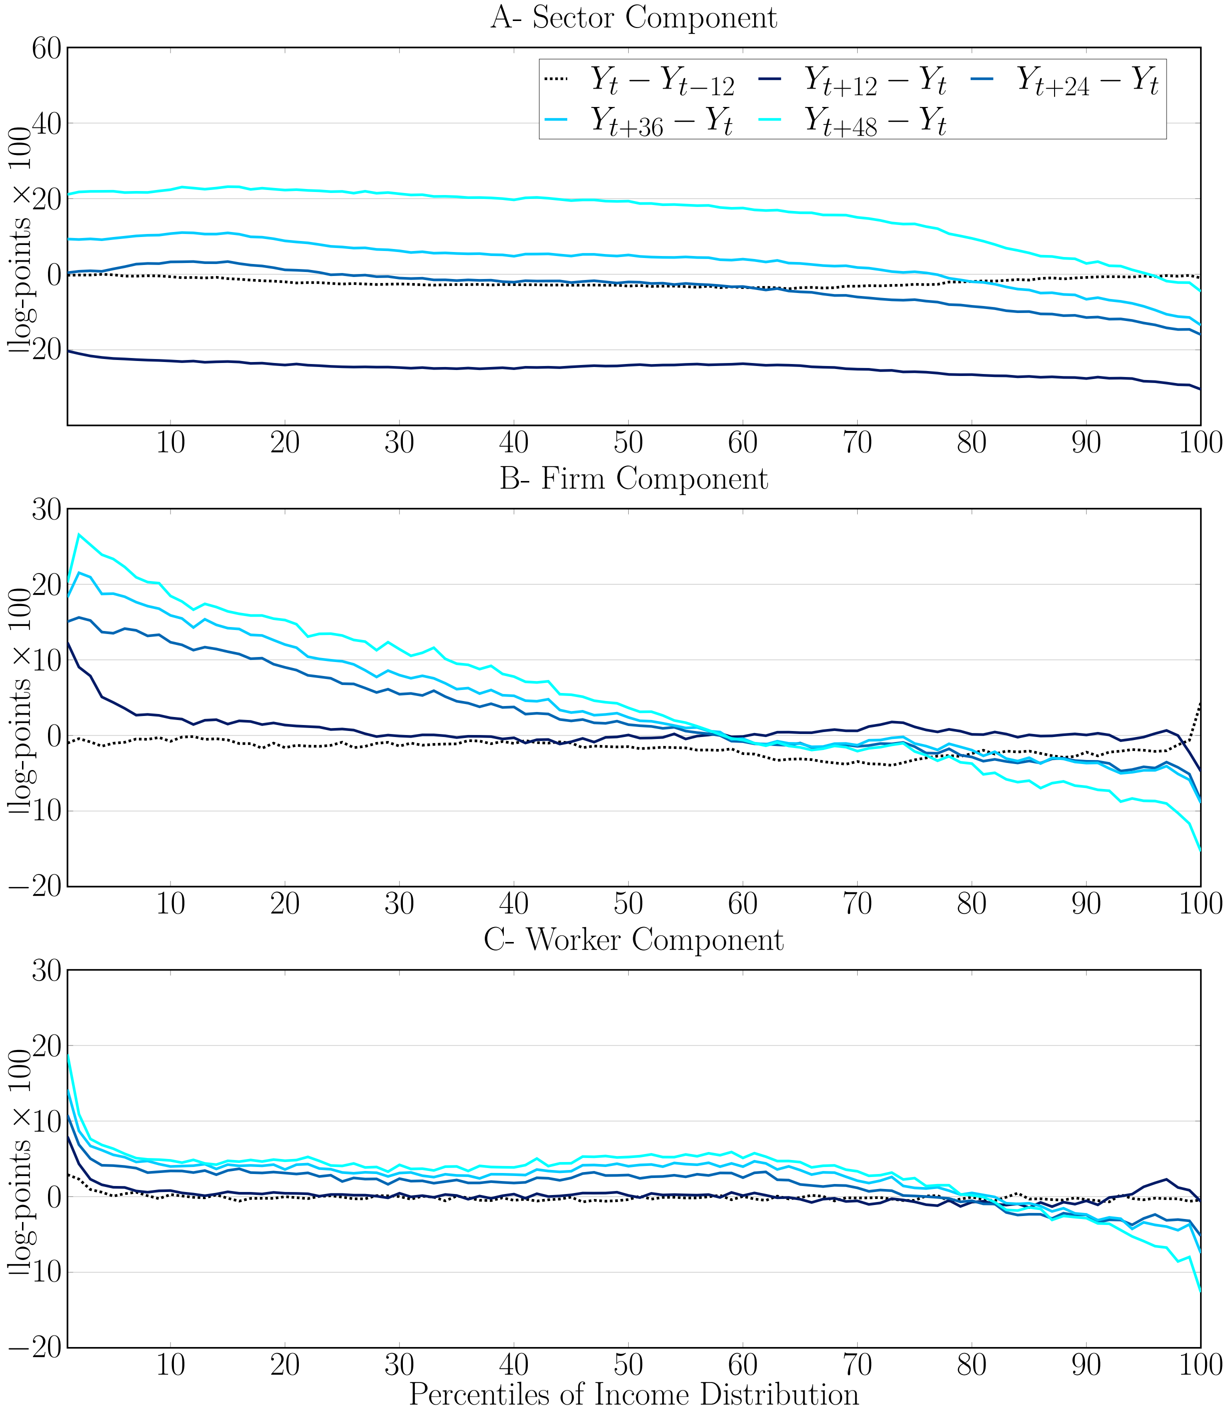

Notes: The figure describes average income growth conditional on the percentile of the distribution of average monthly real income during 2000-2001. The sample is restricted to workers who had at least 6 months of employment during the 2000-2001 period. Panel A replaces a worker's labor income with the average labor income in the sector of employment. Panel B replaces a worker's labor income with the average labor income in the firm of employment net of the sectoral average labor income. Panel C replaces a worker's labor income with the worker's labor income net of the firm's average labor income.

during devaluations? Several papers have documented differences in mobility patterns across groups of workers. For example, Karahan *et al.* (2019) show that in the US: i) the number of employers during the working life (resp. the fraction of job stayers) is decreasing (resp. increasing) in lifetime earnings, and ii) the separation and job-to-job transition rates are declining in workers' earnings. We provide an answer to our question in two steps. First, we document the incidence of different types of transitions and the conditional average income growth by type of transition across the income distribution. Second, we compute a set of counterfactual income dynamics without income changes after different types of transitions.

Workers at the bottom (resp. top) of the pre-devaluation income distribution experienced separation shocks at a higher (resp. lower) rate, but on average their income increased (decreased resp.) with each transition. Figure 8-Panel A plots the cumulative probability of experiencing a separation over the first 4 years after the devaluation as a function of a worker's pre-devaluation income (the same ranking of workers as in Figure 6). This probability is monotonically decreasing in the position of the distribution, with the exception of workers above the 90th percentile. Relatedly, Figure 8-Panel B plots the average income growth across all job transitions that involve an unemployment spell within percentiles of the distribution. In the first year after the devaluation, workers below the 40th percentile experienced an average income growth of 10.4%, while workers above the 40th percentile experienced an average growth of -26%.[11] Four years after the devaluation, workers below the 50th percentile experienced an average growth of 4.3% during job changes that involved an unemployment spell, while the losses of high-income workers were smaller (-9% on average).

Low-income workers were also more likely to make job-to-job transitions and to experience a larger income growth on average when making such transitions. Figure 8-Panels C and D plot the same objects for the case of job-to-job transitions. Qualitatively, the patterns are the same as those observed for separations. The only difference is that, starting from the second year after the devaluation, workers in all percentiles experienced a positive income growth after a job-to-job transition on average. Importantly, the average income growth is still decreasing in the position in the income distribution.

We document that labor mobility is an important driver of the heterogeneous recovery of income. For this, we construct several counterfactual series of income. First, we compute counterfactual income dynamics without changes due to job-to-job transitions. For each worker, we compute changes in income for each pair of subsequent observations over time. Then, we identify the changes in income that are due to job-to-job transitions and replace them with a zero.[12] Finally, we reconstruct for each worker the time series of the level of labor

[11] The labor income drop after a separation is consistent with Davis and Wachter (2011).

[12] Formally, let $Y^i_{t(j)}$ be the j-th chronological observation of worker i in period $t(j)$ in the dataset. Then, we can write

$$Y^i_{t(j)} = Y^i_{t(1)} + \sum_{l=2}^{j} \Delta Y^i_{t(l)},$$

Figure 8 – Income Mobility across the Income Distribution

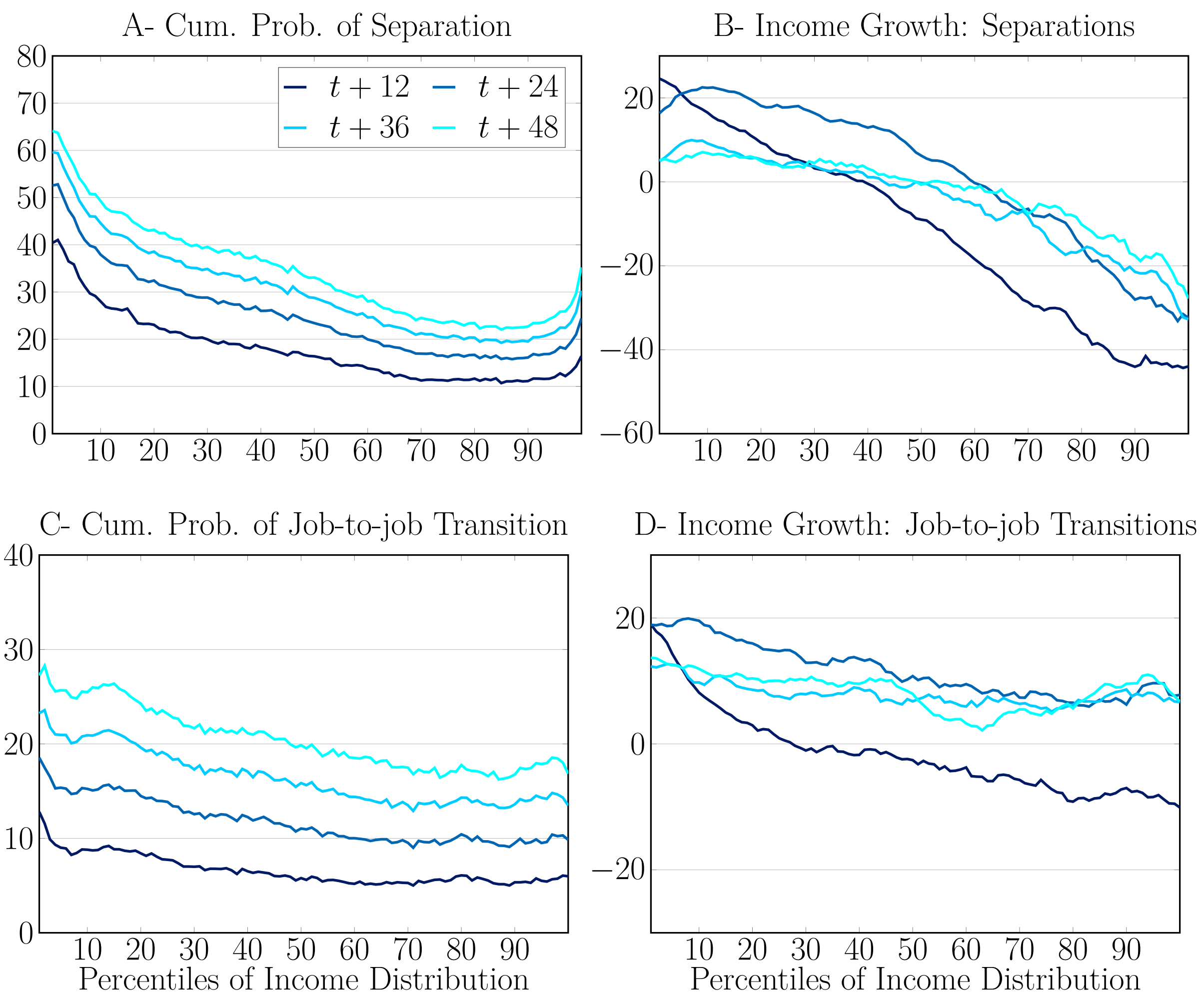

Notes: Panel A plots the cumulative probability of experiencing a separation between December 2001 and December in the next 4 years. Panel B plots the average difference between the (log) income in the new job found after a separation during each year after the devaluation and the (log) income in the previous job. Panel C plots the cumulative probability of experiencing a job-to-job transition between December 2001 and December in the next 4 years. Panel D plots the average difference between the (log) income in the new job found after a job-to-job transition during each year after the devaluation and the (log) income in the previous job. All figures are conditional on the percentile of the distribution of average monthly real income during 2000-2001. The sample is restricted to workers that had at least 6 months of employment during the 2000-2001 period. We truncate the distribution of income changes by the 1% and 99% percentiles to construct Panels B and D.

income with these counterfactual income changes. These counterfactual income dynamics reflect the actual income growth for incumbent and separating workers, and omit income growth experienced during job-to-job transitions.

where $\Delta Y^i_{t(l)} \equiv Y^i_{t(l)} - Y^i_{t(l-1)}$. Then, we construct a counterfactual series for $Y^i_{t(j)}$ by setting $\Delta Y^i_{t(l)} = 0$, whenever the worker i makes a job-to-job transition between $t(l-1)$ and $t(l)$ (i.e, whenever employers differ in those two periods and $t(l) - t(l-1) \leq 1$).

Figure 9 – Contrafactual Income Growth across the Distribution

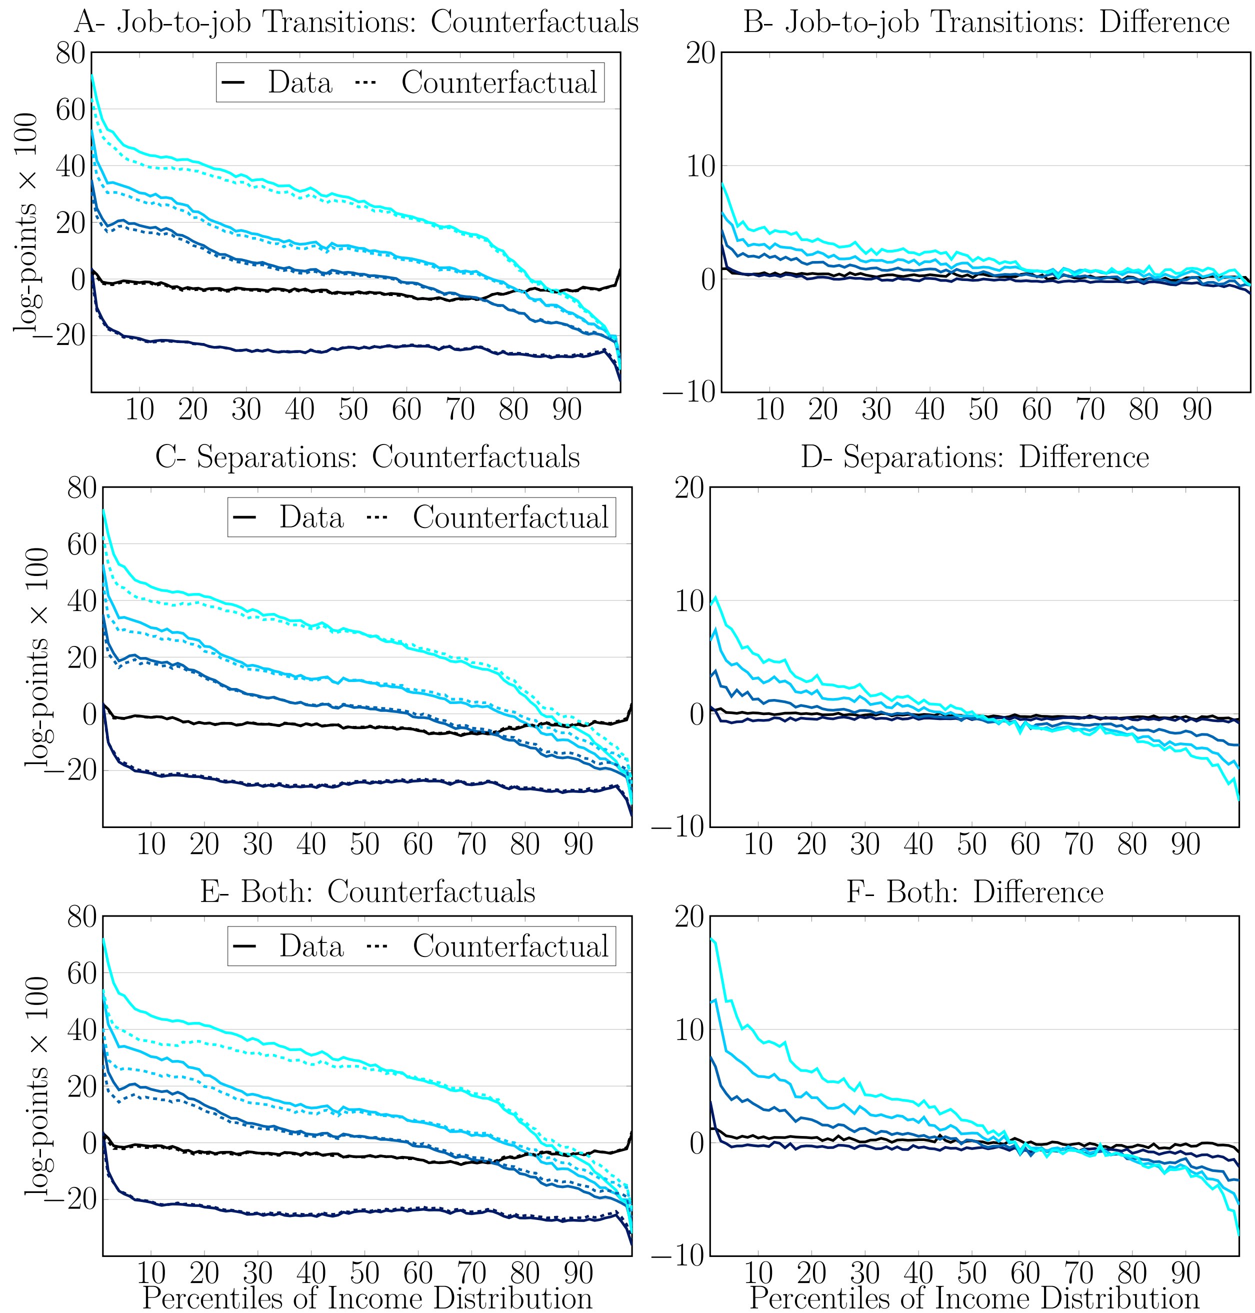

Notes: Panel A describes both the actual average income growth and the counterfactual income growth that omits income changes experienced during job-to-job transitions. Panel B plots the difference between the actual and the counterfactual dynamics to ease the comparison. Panel C describes both the actual average income growth and the counterfactual income growth that omits income changes experienced after separations. Panel D plots the difference between the actual and the counterfactual dynamics to ease the comparison. Panels E and F present similar results for the combined effects of job-to-job transitions and separations. All figures are conditional on the percentile of the distribution of average monthly real income during 2000-2001. The sample is restricted to workers that had at least 6 months of employment during the 2000-2001 period.

Figure 9-Panel A compares the baseline results with the counterfactual income dynamics (for ease of exposition, Figure 9-Panel B plots the difference between both lines). We can see that job-to-job transitions did not generate any heterogeneous income growth before or immediately after the devaluation. However, during the recovery phase, we see that job-to-job transitions positively contributed to higher income growth, especially for workers below the 50th percentile. Quantitatively, job-to-job transitions generated a significant fraction of the pivoting observed in Panel B of Figure 7 (which also shows similar changes before and immediately after the devaluation, followed by positive income growth for workers below the 50th percentile).

Next, we perform a similar exercise with the aim of quantifying the role of mobility due to separations. In this case, we identify the changes in income that are due to separations and replace them with a zero. With these counterfactual growth rates, we reconstruct the time series of labor income for each worker. Figure 9-Panel B shows the results. In this case, the pivoting that can be attributed to income growth generated by separations is even stronger.

Combining both results, labor mobility generates 44% of the pivoting effect at the firm level. Panels E and F combine the effects of both types of labor mobility. The average cumulative income growth for workers at the 10th percentile was 9.2%. Instead, workers at the 90th percentile experienced an average cumulative income growth of -2.2%. As a comparison, the average cumulative income growth for workers at the 10th (resp. 90th) percentile in Panel B of Figure 7 was 18.5% (resp. -7.2%).

As highlighted by Postel-Vinay and Robin (2002), in a model with on-the-job-search, workers are able to extract rents from their current employers when receiving external offers without having to actually move to a new employer. In our analysis, we quantify the effects of labor mobility by exploring the income dynamics of workers who had the opportunity to move to another job, and actually moved. Therefore, our counterfactual analysis only provides a lower bound for the role played by labor mobility.

5.4 Economic Mechanism II: Heterogeneous Income Floors

Given the importance of between-sector and between-workers components for explaining income growth above the median, it is natural to study the heterogeneous dynamics of sector-by-occupation income floors set by unions. Income floors are common labor market institutions across countries. For example, in the U.S. there are state-level minimum wages.[13] While a minimum wage exists in Argentina, it only applied to 1% of workers before the 2002 devaluation. Instead, the main income floor in Argentina is set by trade unions and

[13] Several papers have studied the macroeconomic consequences of a minimum wage. See, e.g., Engbom and Moser (2018) for a study of minimum wages in Brazil, Flinn (2006) in the US, and Harasztosi and Lindner (2019) in Hungary.

they differ by sectors and occupations, as is typically the case in many European countries. Can unionization status explain the heterogeneous individual recoveries across the income distribution?

The answer to this question is yes, which we explain in three steps. First, we briefly describe the role of unions in Argentina. Second, we present evidence on the role of unions for income growth within sectors.[14] Finally, we reproduce our main fact by unionization status and find significant differences.

In Argentina, a single union has monopoly power to represent workers and negotiate a CBA at a sectoral level. A CBA determines the minimum labor income for all workers in that sector employed in a subset of occupations, regardless of their individual membership status. By law, negotiated wages must be above the national minimum wage. For the largest firms in a sector, unions also negotiate firm–specific CBAs, which have to offer better terms to workers than the sectoral CBA.

Following the 2002 devaluation, unionized workers whose incomes were covered by a CBA saw their labor income recover faster than non-unionized workers. Figure 10 plots income by unionization status over time for some sectors with strong unions. The figure plots the average income in the CBA across occupations and the average income of workers covered and non-covered by the CBA. Covered workers are those who are unionized according to the SIPA dataset, and whose labor income is within the prevailing range of incomes established by the CBA in October 2002. We choose October 2002 since between 1995 and that date unions did not renegotiate their CBAs. By law, an expired CBA still remains legally binding until a new one is negotiated.[15]

For all the industries in which bargained income changes were above inflation, the average nominal income growth of covered workers was 30% higher than non-covered workers. The labor income growth rate of unionized workers closely follows the average growth specified by the CBA. This pattern holds for retail trade, construction, motor vehicle manufacturing, and freight transport by road. Finally, there is one sector (mechanics) in which the CBA's income growth is almost equal to the cumulative inflation between 2002 and 2005. In that sector, income growth between 2002 and 2005 does not vary by unionization status. We conclude that there is a significant heterogeneity of income growth by unionization status in sectors with strong unions. Importantly, the timing of union renegotiations coincides with the moment when income inequality began to decline (2nd semester of 2003).

Until now, we illustrated the role of unions in a subset of sectors. Next, we present the

[14] As we showed above, the primary source of heterogeneous recovery of labor income is within sectors.

[15] These groups are constructed using the unionization status variable in the SIPA data. The unionization variable becomes available in June 2003, and presents a high degree of persistence in the sample. For this reason, we are confident that the majority of these workers maintained their unionization status between October 2002 and June 2003. Since unions negotiate a minimum monthly labor income for specific occupations, we added the second condition to identify workers near the prevailing minimum income in October 2002.

Figure 10 – Normalized Labor Income by Union Coverage and Labor Income in CBAs

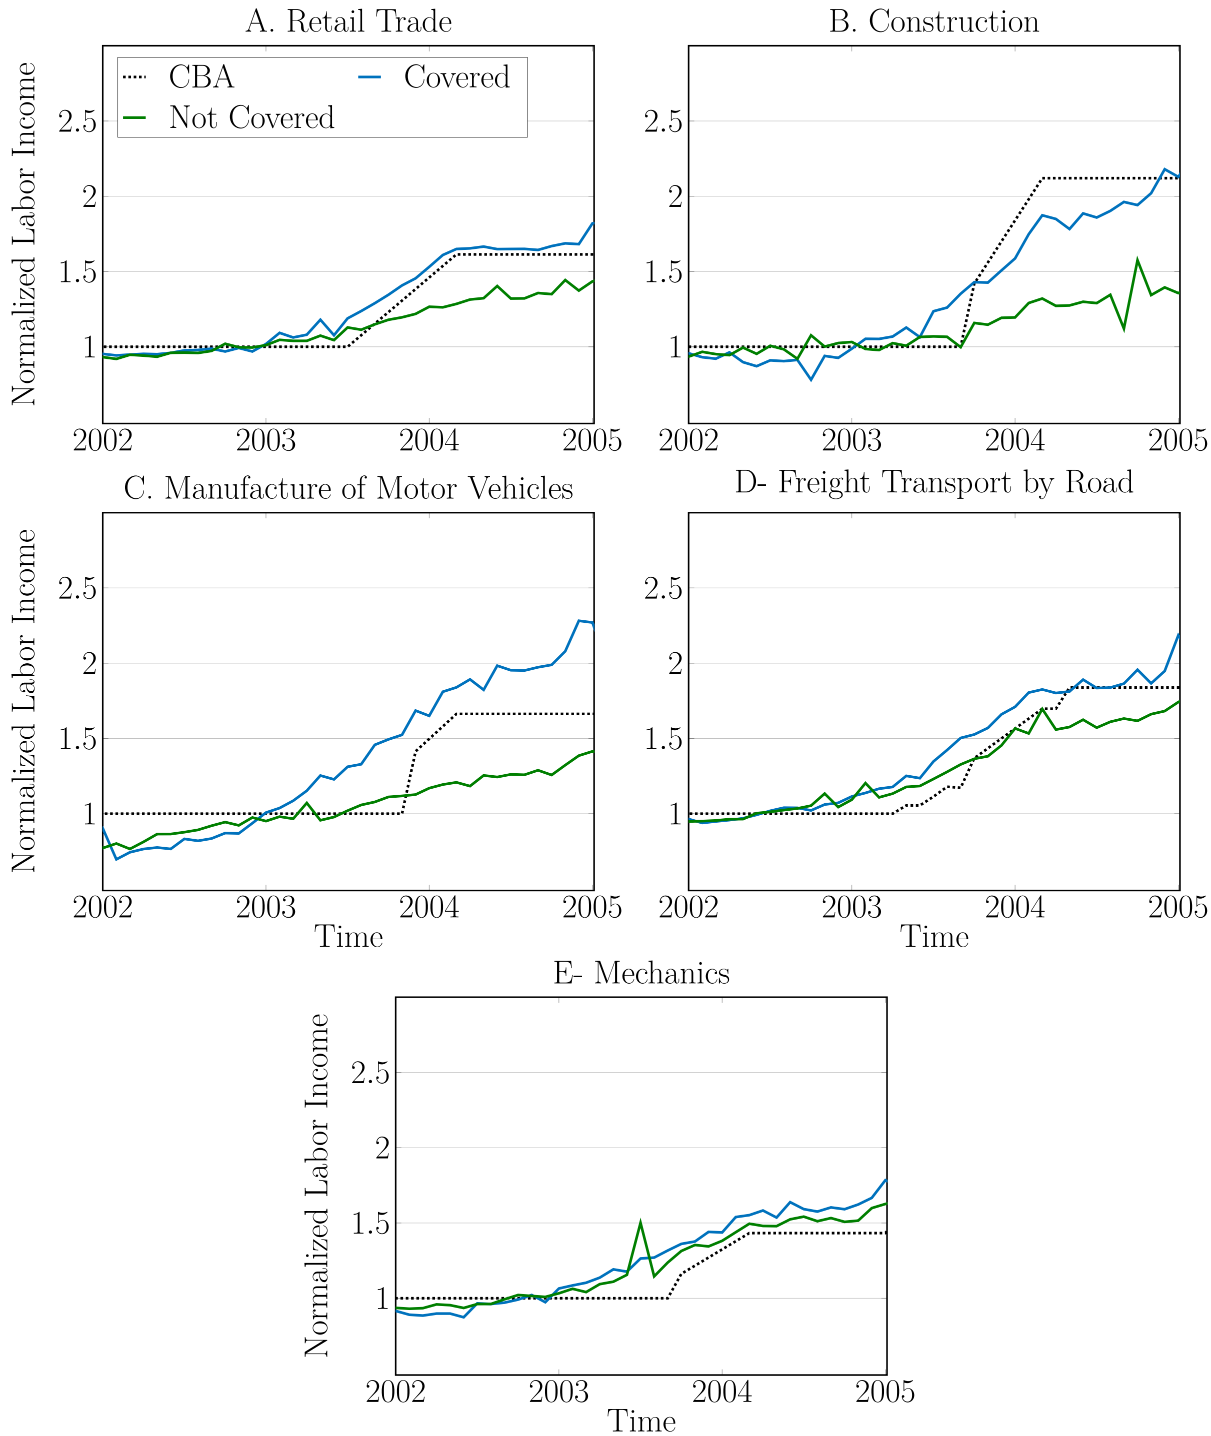

Notes: Panels A to E plot average nominal income in the CBA across occupations and the average nominal income of workers covered and non-covered by the CBA across five industries (i.e.,). A worker belongs to the group "Covered" if it is unionized in June 2003 according to SIPA dataset and her labor income is between the lowest and highest incomes across occupations in CBAs. A worker belongs to the group "Not Covered" if it is not unionized in June 2003 in SIPA dataset.

Figure 11 – Average Income Growth Conditional on Average Income 2000-2001 by Unionization Status

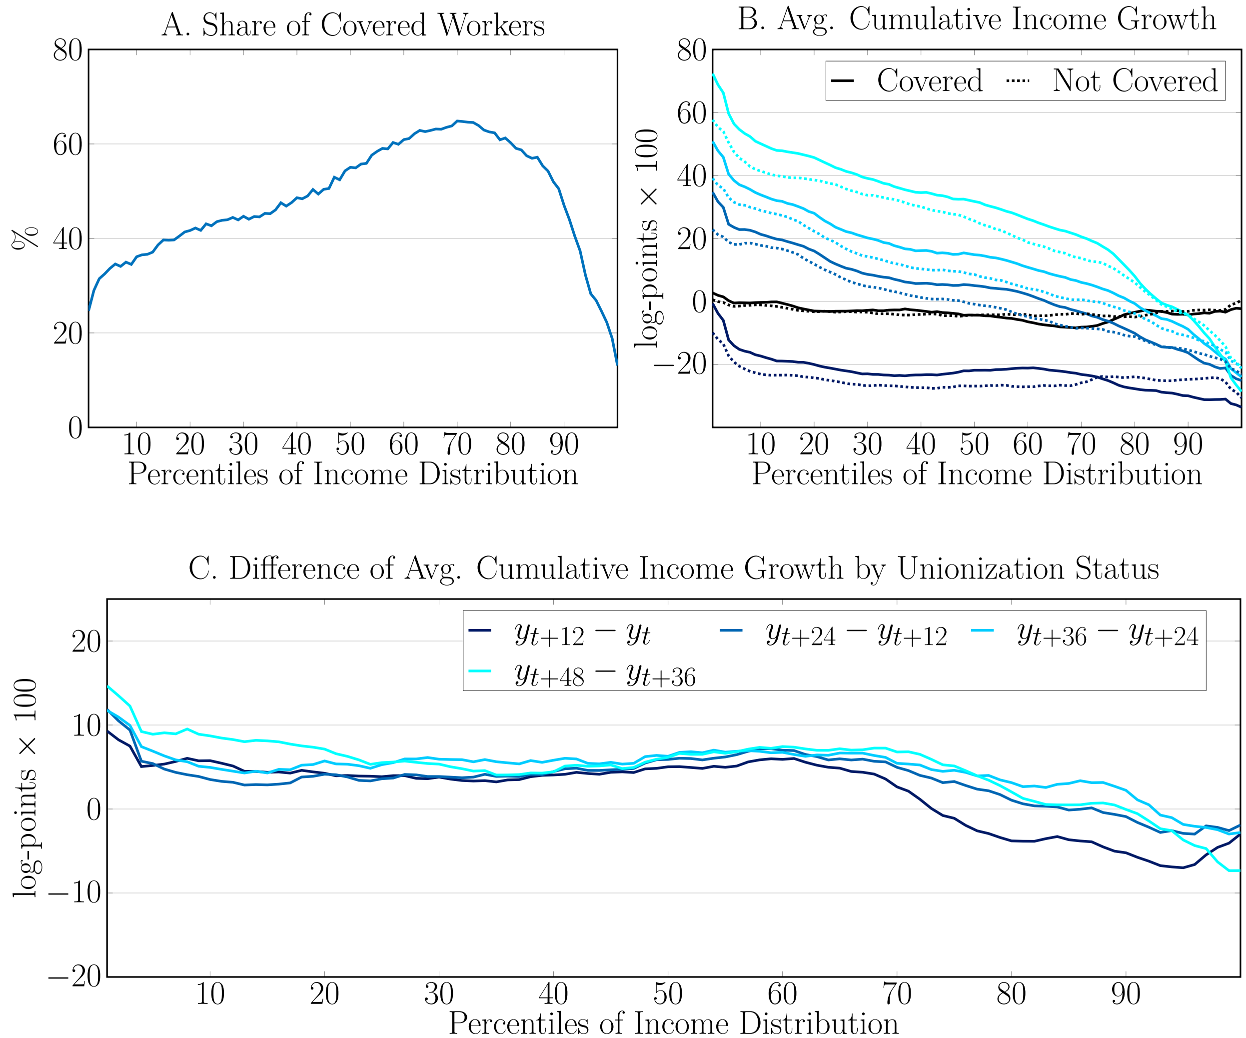

Notes: Panel A shows the shared of unionized workers by percentiles of income, as in Figure 6. Panel B shows average cumulative income growth by percentiles and unionization status. Panel C shows the difference in the average cumulative income growth between unionized and not unionized workers by percentiles.

contribution of unionization status to the main fact of this paper. Figure 11 reports the share of unionized workers and average labor income growth by unionization status as a function of pre-devaluation income. To construct Figure 11 we split workers according to their unionization status only, regardless of their incomes relative to bargained incomes.[16]

The share of unionized workers is increasing for the worker between the 1st and 70th percentiles and decreasing from the 70th to 100th percentile. The share of unionized workers

[16]The digitalization of all industries' CBAs and the merge with SIPA data is outside the scope of this paper. Each industry has its own industry specific contract format that changes over time. Therefore, we were not able to standardize CBAs across all industries. Nevertheless, we reproduce Figure 10 in the Online Appendix with different definitions of coverage to show how these definitions affect the measurement of income dynamics by unionization status (see Figure C.13).

is above 40% between the 20th and 80th percentiles. Thus, union bargaining is relevant primarily for workers in the middle- to top- of the income distribution, and less so for workers at the top and the bottom of the income distribution.

The average cumulative income growth was higher for unionized workers than non-unionized workers. The average difference in income growth across unionization status is close to zero between December 2002 and December 2001. If we focus on workers below the 70th percentile, the differences in income growth for unionized workers relative to non-unionized workers increased over time. The average difference one year after the devaluation was 4%, and that difference became 6% four years after the devaluation. On the other hand, unionized workers at the top of the pre-devaluation income distribution experienced a slower recovery relative to non-unionized workers (see Card, 1996, for similar evidence of a smaller/negative union premium at the top of the distribution in the US). The difference in relative income growth for workers below and above the 70th percentile is quantitatively important, as it resembles the pattern of heterogeneous recovery of the workers' components of income presented in Panel C of Figure 7.

6 Additional Mechanisms and Robustness

6.1 Additional Mechanisms

Sectoral trade exposure. Can trade exposure explain the heterogeneous individual recoveries across the income distribution? The answer is no. Sectoral trade exposure correlates strongly with labor income growth, but it cannot generate a decline in inequality. Here, we demonstrate this result by focusing on a broad classification of trade exposure, i.e., tradable and nontradable sectors. We relegate our analysis of sectoral trade exposure at the three-digit SIC level to Section D.1 in the Online Appendix.

For trade exposure to explain the heterogeneous individual recoveries across the income distribution, two conditions need to be satisfied. First, the NER and sectoral labor income must be correlated, and their correlation should be a function of trade exposure. For example, workers in the nontradable sector should benefit the least from a devaluation. Second, workers' pre-devaluation permanent income and the income recovery in sectors that benefit the most from a devaluation must be negative correlated. For example, if the non-tradable sector recovers more slowly, workers at the top of the income distribution should concentrate in this sector. While we find strong support for the first condition, we do not find support for the second condition.

Figure 12-Panel A (resp. B) shows sectoral income growth in the tradable (resp. non-tradable) sector. Figure 12-Panel C shows the difference of sectoral income growth across tradable and non-tradable sectors. We construct these figures in two steps. First, we group workers according to their position in the pre-devaluation income distribution, as in Sec-

tion 5, and their sector of employment (i.e., tradable or nontradable).[17] Second, for each percentile of the income distribution and broad sector (i.e., tradable and nontradable), we compute the average income growth combining sectoral income growth at the four-digit SIC level and the composition of workers across those sectors.

Figure 12 – Average Income Growth Conditional on Average Income 2000-2001 for Tradable and Nontradable Sectors

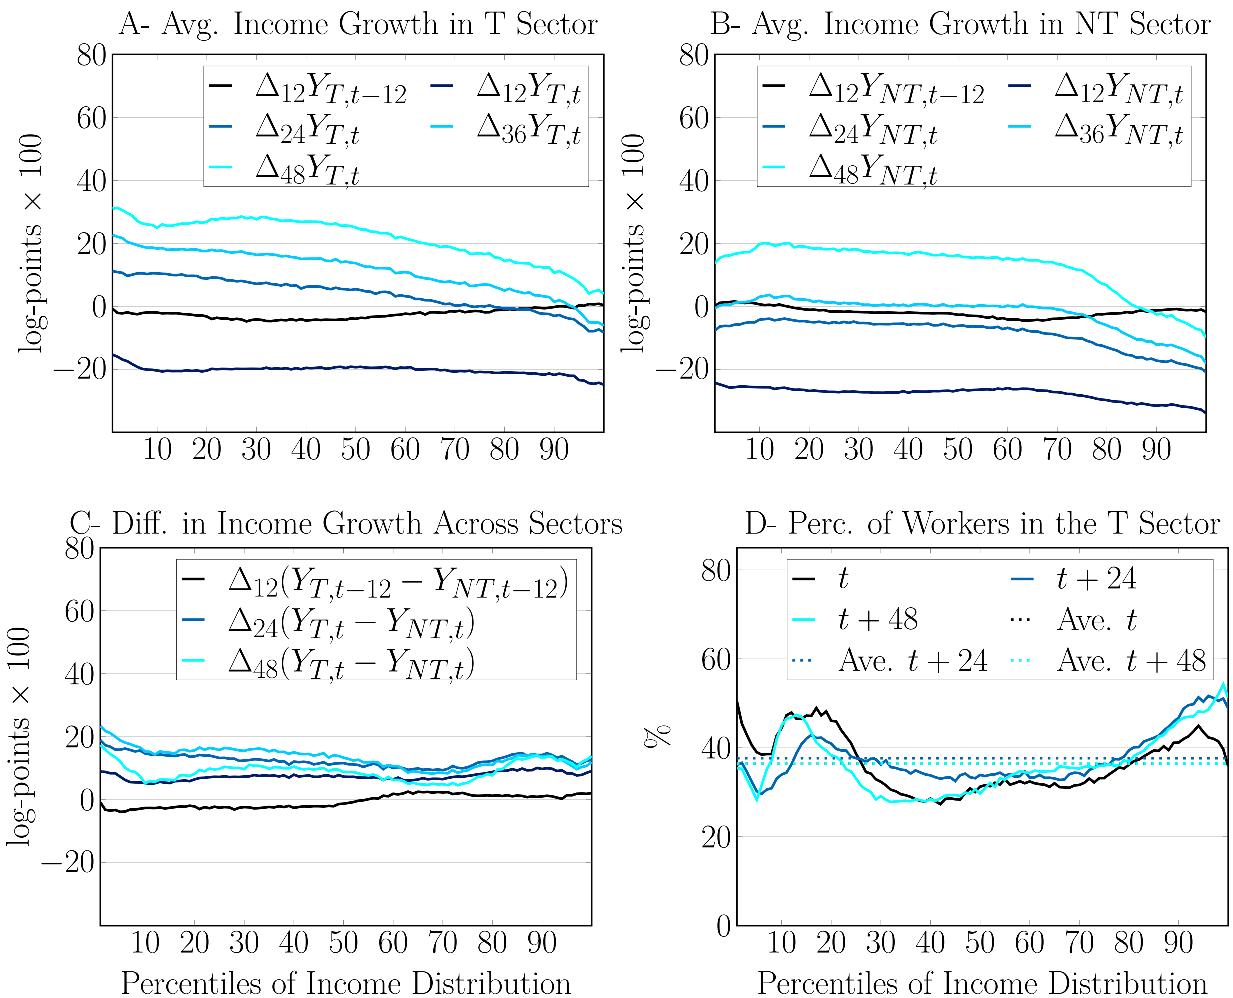

Notes: Panel A (resp. Panel B) plots average income growth in the tradable (resp. nontradable) sector conditional on the percentile of the distribution of average monthly real income during 2000-2001. Panel C plots the difference between average income growth in the tradable and nontradable sectors. Panel D plots the percentage of workers in the tradable sector in December 2001 ("t"), in December 2003 ("t+24"), and in December 2005 ("t+48"), together with the average across the percentiles of the income distribution in those dates. The sample is restricted to workers that had at least 6 months of employment during the 2000-2001 period.

The labor income of tradable-sector workers permanently increased by 10% relative to that of nontradable-sector workers after the devaluation. The average difference across per-

[17] The tradable sector includes agriculture, livestock, and hunting, fishing and related services, mining, and the manufacturing industry.

centiles of income is 0.5% in favor of the nontradable sector between December 2000 and 2001.[18] Following the 2002 devaluation, there is a faster recovery of tradable-sector labor income than nontradable sector labor income. The average differences across percentiles over the course of 4 years are 7%, 12%, 13%, and 9% in chronological order. In conclusion, there is a significant difference in labor income dynamics across the tradable and nontradable sectors resulting from the predicted increase of revenue in tradable sector relative to the non-tradable sectors.

There is no quantitatively relevant "pivoting" pattern in the share of tradable workers across the pre-devaluation income distribution. That is, the share of tradable employment is not decreasing in pre-devaluation income. Figure 12-Panel D plots the share of tradable workers in December 2001 and two and four years after. The larger share of tradable workers across the distribution is in the 15th percentile, and it is 10% higher than the mean across the income distribution. Since the relative income growth of the tradable relative to the non-tradable is 10%, the exposure to the tradable sector can only explain at most 1% difference across sectors. Furthermore, there is a concentration of tradable employment at the top of the distribution, which goes against our main finding.

Finally, we find a small reallocation of labor towards the tradable sector after the 2002 devaluation. As Figure 12-Panel D shows, the average share in the tradable sector was 36% in December 2001 and increased to 37% in December 2004. In light of this large devaluation, the lack of mobility into the tradable sector is surprising. Nevertheless, it is consistent with the fact that the tradable premium persists for such a long time.

Changes in labor income risk. Can the decrease in inequality be explained by a lower labor income risk? To illustrate the logic of this question, suppose that the income process follows a standard AR(1) process. Then, a decrease in the standard deviation of the innovation would translate into a compression of the stationary distribution. Thus a decrease in the standard deviation of income growth could explain a lower level of inequality.

One potential source of a decline in labor income risk is the observed sharp decrease in the separation rate after the 2002 devaluation (results available upon request). The literature has previously documented that job displacements are typically associated with large cumulative earnings losses (see, for example, Davis and Wachter, 2011). Thus, if the incidence of such large negative events decreases, the distribution after the devaluation could become more equal.

Nevertheless, the requirements for this mechanism to work are not observed in the data: Inequality decreased *despite* an increase in the standard deviation of income growth. Figure 13 shows selected moments of the labor income growth distribution. During the recession and before the devaluation, the interquartile range of the distribution of labor income growth was

[18] Figure D.3 in the Online Appendix shows that the relative wages of tradable workers was slightly decreasing relative to non-tradable workers between 1994 and 2001.

Figure 13 – Moments of the Distribution of Labor Income Growth

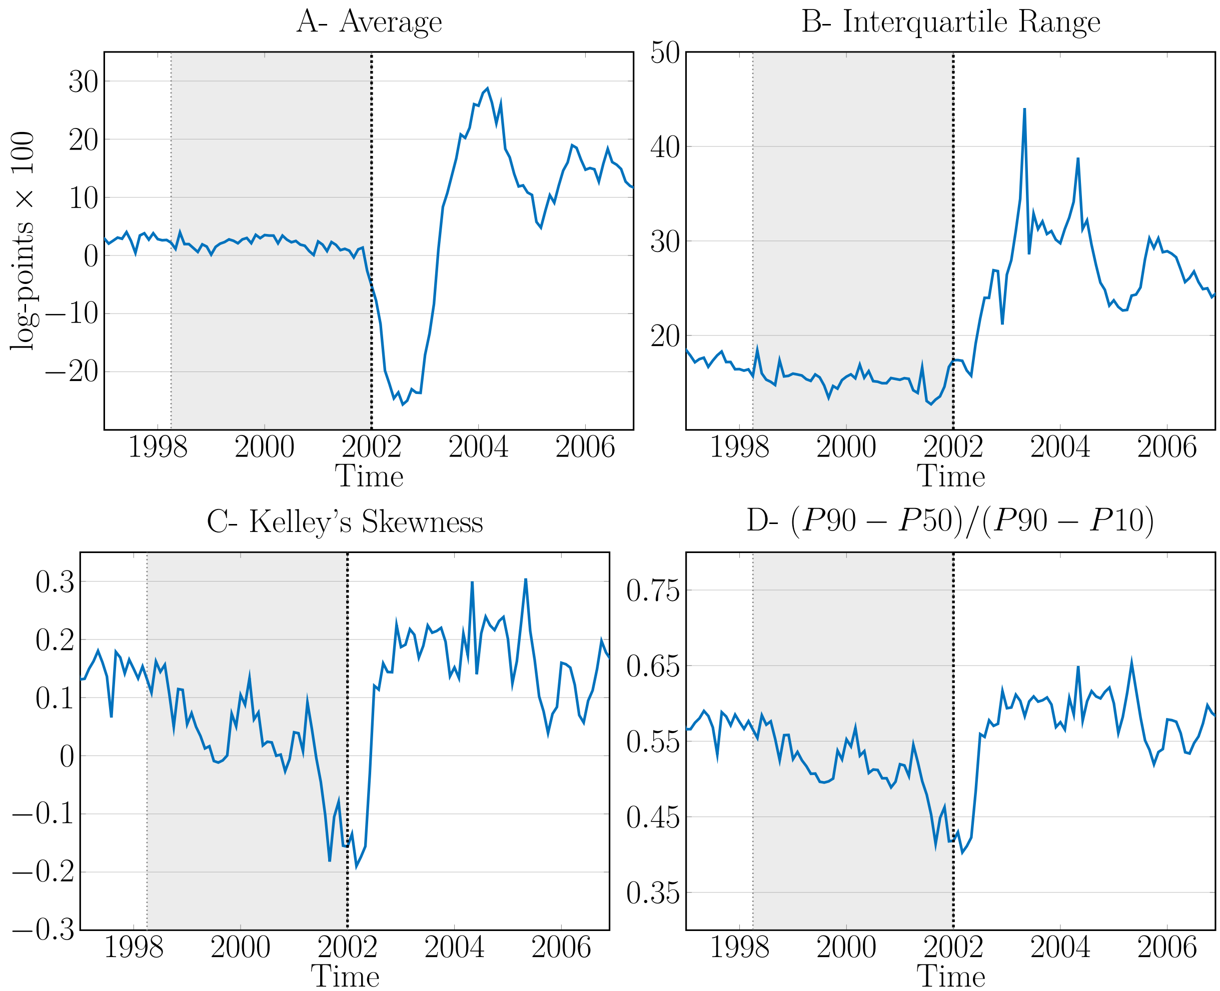

Notes: Panels A to D plot (in the following order) the average, the interquartile range, Kelley's skewness ($\frac{(P90-P50)-(P50-P10)}{P90-P10}$), and the decomposition of the Kelley's skewness ($\frac{P90-P50}{P90-P10}$) of year-over-year income growth from 1997 to 2007.

almost constant, and Kelley's skewness continuously decreased (similar patterns have been documented for the U.S. by Guvenen, Ozkan and Song, 2014).[19] After the devaluation, there was a significant increase in the dispersion of year-over-year income growth. Figure 13-Panel B shows a sharp and persistent increase in the interquartile range of year-over-year income growth from below 20% up to 40%. Moreover, the increase in dispersion is not symmetric. After the devaluation, there was a reversal in the negative trend in the skewness, which changes from -0.2 to an average of 0.15. In other words, the right tail of the distribution of income growth expands. As Panel D shows, most of the movements in skewness come from changes in the distribution above the median: 60% of Kelley's skewness can be attributed to the upper tail after the devaluation.

[19]Kelley's measure of skewness is defined as $\frac{(P90-P50)-(P50-P10)}{P90-P10}$. Since it is based on percentiles, it is more robust to outliers.

Two mechanisms could explain the increase in labor income risk. First, a larger reallocation of labor, since the reallocation of workers across employment states, firms, and sectors is associated with large income variations, as previously shown. The fact that the standard deviation of income growth of job stayers also increases (see Figure D.6 in the Online Appendix) points to an additional mechanism: the heterogeneous arrival rates of adjustment times of nominal income in the short run after a devaluation and heterogeneous growth in real income conditional on adjustment in the medium run.

6.2 Robustness

This section analyzes the role of policy changes or additional dimensions of the labor market to better understand the heterogeneous labor market dynamics during devaluations. We provide a summary of the results here, and the complete analysis can be found in Online Appendix Sections D.3-D.6.

Changes in the minimum wage. The nominal monthly minimum wage in Argentina was fixed at $200 from August 1993 to July 2003.[20] After the 2002 devaluation, there was a continuous drop in the real minimum wage until its first adjustment in July 2003. Since then, it has experienced a series of increases, and by the end of 2005 its real value was equivalent to the 10th percentile of the real income distribution.

We provide evidence showing that changes in the real minimum wage could not have been the main driver behind the post-devaluation drop in inequality. First, we show that the timing of this potential explanation is misaligned. Six months after the devaluation, divergent dynamics of the bottom and top percentiles of the income distribution emerged. This occurred while the real minimum wage kept *decreasing* due to a lack of adjustment and became even less binding. Thus, the drop in inequality preceded the increase in the real minimum wage. In addition, it is worth pointing out that after the large increase in the real minimum wage in September 2004, of more than 20 log points, we do not see any further large changes in inequality.

Second, the heterogeneous recovery we observe in Figure 6 is almost a linear function of the position of a worker in the pre-devaluation income distribution. It is highly unlikely that changes in the minimum wage had spilled over up to the 80th or 90th percentile in such a short period of time.

Changes in hours versus hourly wages. Throughout the paper, we report facts about monthly real labor income and not hourly wages, due to data limitations. Nevertheless, we

[20] Given the lack of adjustment for such a long period of time, it is not surprising that the monthly minimum wage became binding for a small fraction of the population. In 2001, it was equivalent to the monthly nominal income of a worker in the 2nd percentile of the income distribution.

performed three exercises to show that the main facts presented in Section 5 are driven by changes in hourly wages and not by fluctuations in hours worked.

In the first exercise, we computed average weekly hours and the distribution of average weekly hours by quintiles of the income distribution in the household survey, which contains information on hours of work. While average hours worked drop by at most 2% after the 2002 devaluation, this magnitude cannot explain the drop in real labor income of almost 30%. We also find almost no difference in the evolution of hours worked across quintiles of the income distribution. Moreover, the small drop in hours is homogeneous across the income distribution. Thus, differences across income groups cannot account for the large decrease in inequality. In the second exercise, we analyzed workers' real hourly wages using the same data. We find that the dynamics of the hourly wage distribution closely follow the dynamics of the monthly income distribution.

In the last exercise, we divide workers according to their full- and part-time status using information on workers' type of labor contract in the SIPA dataset.[21] We find quite similar dynamics of the mean real labor income across groups of full- and part-time workers. We also find similar dynamics of the interquartile range and the standard deviation of the labor income distribution across groups.

Worker-specific inflation. If we are interested in a worker's consumption possibilities after a devaluation, the appropriate deflator for a worker's nominal income ($\tilde{y}_i$) should be based on each worker's consumption basket (p_i) instead of the aggregate CPI (p). We can decompose the measure of real income of interest as $\tilde{y}_i/p_i \equiv (\tilde{y}_i/p) \times (p/p_i)$. In this paper, we focus on the first component, $\tilde{y}_i/p$, to render the comparison of real income dynamics across workers more transparent.

In Online Appendix Section D.5, we reproduce Figure 6 by constructing measures of income-specific deflators. Using micro-data from the national expenditure survey in Argentina, we document that it is indeed the case that households with lower incomes experienced a higher inflation rate after the devaluation, since their consumption basket is tilted toward goods with prices that comove more with the nominal exchange rate (as in Cravino and Levchenko, 2017). However, these differences in income-specific inflation rates are not large enough to overturn our main fact.

The informal labor market. In Online Appendix Section D.6, we provide a broader picture of the Argentinian labor market during devaluations by extending the analysis to the informal labor market. First, we find that labor income also decreased in the informal

[21] The full-time group includes workers with and without a termination date specified in their contracts. The part-time group also includes seasonal workers, trainees, and temporary workers. In order to be overly cautious, we also include in this group all workers in the agriculture, mining, fishing, and construction sectors due to the intermittent working periods common in these sectors.

market after devaluations. In fact, the drop is larger and even more persistent that in the formal sector. However, we do not see a clear compression of the cross-sectional distribution of informal income. This is consistent with the fact that unions—which are present only in the formal sector—explain a faster recovery of real incomes.

When we examine the dynamics of informal employment, we find that the share of informal employment decreases after the devaluation, which is in line with improving conditions in the formal labor market. As the decline of the informality rate is associated with transitions from the informal to the formal sector (which on average pays higher wages), this finding suggests that labor mobility played an additional role in compressing the overall income distribution.

7 Discussion: Linking Empirical Evidence and Theory

Our measurement exercise points to important considerations (i.e., the role of firms vis-a-vis sectors and occupations) and general mechanisms (i.e., labor mobility and income floor dynamics) through which economies adjust after large NER devaluations.

While our findings about inequality might be of interest on their own, our results also provide empirical guidance to some prominent theories. Following the seminal paper of Bils and Klenow (2004), there has been far greater theoretical work on price dynamics relative to wage dynamics, despite the importance of wage rigidities in macroeconomic fluctuations (see Christiano, Eichenbaum and Evans, 2005). We conjecture that this differential development occurred in response to the availability of high-quality pricing data and the relative dearth of high-quality wage data. This paper provides an empirical footing for new theoretical work on the macroeconomic consequences of wage rigidities. We show that real labor income decreases during large devaluations due to prices adjusting faster than wages and the the adjustment of wages is heterogeneous across workers.

Our main facts documents the dynamics of inequality in the aftermath of large NER devaluations. Inequality matters for aggregate demand whenever markets are incomplete. The implications of inequality has been studied for monetary policy shocks (see, e.g., Auclert, 2019, Kaplan, Moll and Violante, 2018, Auclert and Rognlie, 2018), but not for sudden stops. Extending canonical models of large devaluations (see, e.g., Mendoza, 2010) could help us understand these events in two dimensions. First, inequality could affect aggregate dynamics. Second, the distributional impact of sudden stops could be a primary consideration for the design of policies with inequality concerns.

References

Ates, S. and Saffie, F. (2016). Fewer but better: Sudden stops, firm entry, and financial selection. *FRB International Finance Discussion Paper*, (1187).

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, **109** (6), 2333–67.

— and Rognlie, M. (2018). *Inequality and aggregate demand*. Tech. rep., National Bureau of Economic Research.

Barattieri, A., Basu, S. and Gottschalk, P. (2014). Some evidence on the importance of sticky wages. *American Economic Journal: Macroeconomics*, **6** (1), 70–101.

Benguria, F., Matsumoto, H. and Saffie, F. (2020). Productivity and trade dynamics in sudden stops.

Bils, M. and Klenow, P. J. (2004). Some evidence on the importance of sticky prices. *Journal of political economy*, **112** (5), 947–985.

Blaum, J. (2019). Global firms in large devaluations. *Unpublished Manuscript, Brown University.*

Buera, F. J. and Kaboski, J. P. (2012). The rise of the service economy. *American Economic Review*, **102** (6), 2540–69.

Burstein, A., Eichenbaum, M. and Rebelo, S. (2005). Large Devaluations and the Real Exchange Rate. *Journal of Political Economy*, **113** (4).

Calvo, G. A., Izquierdo, A. and Talvi, E. (2006). Sudden stops and phoenix miracles in emerging markets. *American Economic Review*, **96** (2), 405–410.

Card, D. (1996). The effect of unions on the structure of wages: A longitudinal analysis. *Econometrica: Journal of the Econometric Society*, pp. 957–979.

Christiano, L., Eichenbaum, M. and Evans, C. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, **113** (1), pp. 1–45.

Cravino, J. and Levchenko, A. A. (2017). The distributional consequences of large devaluations. *American Economic Review*, **107** (11), 3477–3509.

Davis, S. J. and Wachter, T. V. (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, pp. 1–72.

Dickens, W. T. *et al.* (2007). How wages change: micro evidence from the international wage flexibility project. *Journal of Economic Perspectives*, **21** (2), 195–214.

Doepke, M. and Schneider, M. (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy*, **114** (6), 1069–1097.

Drenik, A., Pereira, G. and Perez, D. J. (2018). Wealth redistribution after exchange rate devaluations. *AEA Papers and Proceedings*, **108**, 552–56.

Eichengreen, B. and Sachs, J. (1985). Exchange Rates and Economic Recovery in the 1930s. *The Journal of Economic History*, **45** (4), 925–946.

Engbom, N. and Moser, C. (2018). Earnings inequality and the minimum wage: Evidence from brazil. *Federal Reserve Bank of Minneapolis-Opportunity and Inclusive Growth Institute Working Paper*, **7**, 18–50.

Fajgelbaum, P. D. and Khandelwal, A. K. (2016). Measuring the unequal gains from trade. *The Quarterly Journal of Economics*, **131** (3), 1113–1180.

Feenstra, R. C., Inklaar, R. and Timmer, M. P. (2015). The next generation of the penn world table. *American economic review*, **105** (10), 3150–82.

Flinn, C. J. (2006). Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica*, **74** (4), 1013–1062.

Frankel, J. A. and Rose, A. K. (1996). Currency crashes in emerging markets: An empirical treatment+. *Journal of International Economics*, **41** (3), 351–66.

Frías, J. A., Kaplan, D. S. and Verhoogen, E. A. (2009). *Exports and wage premia: Evidence from Mexican employer-employee data*. Tech. rep.

Gasparini, L. and Tornarolli, L. (2009). Labor informality in latin america and the caribbean: Patterns and trends from household survey microdata. *Revista Desarrollo y Sociedad*, (63), 13–80.

Gopinath, G. and Neiman, B. (2014). Trade adjustment and productivity in large crises. *The American Economic Review*, **104** (3), 793–831.

Grigsby, J., Hurst, E. and Yildirmaz, A. (2019). *Aggregate nominal wage adjustments: New evidence from administrative payroll data*. Tech. rep., National Bureau of Economic Research.

Guvenen, F., Ozkan, S. and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, **122** (3), 621–660.

Harasztosi, P. and Lindner, A. (2019). Who pays for the minimum wage? *American Economic Review*, **109** (8), 2693–2727.

Hausman, J. K., Rhode, P. W. and Wieland, J. F. (2019). Recovery from the great depression: The farm channel in spring 1933. *American Economic Review*, **109** (2), 427–72.

Kahn, S. (1997). Evidence of nominal wage stickiness from microdata. *The American Economic Review*, **87** (5), 993–1008.

Kaplan, G., Moll, B. and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, **108** (3), 697–743.

Karahan, F., Ozkan, S. and Song, J. (2019). Anatomy of lifetime earnings inequality: Heterogeneity in job ladder risk vs. human capital. *FRB of New York Staff Report*, (908).

Laeven, L. and Valencia, F. (2012). Systemic banking crises database: An update.

Laeven, M. L. and Valencia, M. F. (2018). *Systemic banking crises revisited*. International Monetary Fund.

Le Bihan, H., Montornès, J. and Heckel, T. (2012). Sticky wages: evidence from quarterly microeconomic data. *American Economic Journal: Macroeconomics*, **4** (3), 1–32.

Mendoza, E. G. (2010). Sudden stops, financial crises, and leverage. *American Economic Review*, **100** (5), 1941–66.

Pinkovskiy, M. and Sala-i Martin, X. (2016). Lights, camera… income! illuminating the national accounts-household surveys debate. *The Quarterly Journal of Economics*, **131** (2), 579–631.

Postel-Vinay, F. and Robin, J.-M. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, **70** (6), 2295–2350.

Schmitt-Grohé, S. and Uribe, M. (2016). Downward nominal wage rigidity, currency pegs, and involuntary unemployment. *Journal of Political Economy*, **124** (5), 1466–1514.

Sigurdsson, J. and Sigurdardottir, R. (2011). *Evidence of Nominal Wage Rigidity and Wage Setting from Icelandic Microdata*. Economics wp55, Department of Economics, Central bank of Iceland.

Song, J., Price, D. J., Guvenen, F., Bloom, N. and Von Wachter, T. (2018). Firming up inequality. *The Quarterly journal of economics*, **134** (1), 1–50.

Uribe, M. and Schmitt-Grohé, S. (2017). *Open economy macroeconomics*. Princeton University Press.

Verner, E. and Gyongyosi, G. (2018). Household debt revaluation and the real economy: Evidence from a foreign currency debt crisis.

Young, A. (2013). Inequality, the urban-rural gap, and migration. *The Quarterly Journal of Economics*, **128** (4), 1727–1785.