Skill Premium, Labor Supply and Changes in the Structure of Wages in Latin America*

Manuel Fernández†

Julián Messina‡

November 2, 2016

Abstract

In Argentina and Chile earnings inequality increased in the 1990s and declined in the 2000s. In Brazil, inequality declined dramatically during the two decades. The paper traces earnings inequality dynamics in Argentina, Brazil and Chile linking the observed trends with changes in the composition of the labor force and the premiums of human capital. The reduction in the education and experience premiums are major drivers of the recent earnings inequality decline. Behind this common result lies an important divergence. The reduction of the experience premium is key to explain falling upper-tail (90/50) inequality. The decline of the education premium has a much larger explanatory power in the reduction of lower-tail (50/10) inequality. We show that the decline of the high-school and college premium is largely driven by the educational upgrading of the workforce, but our estimates suggest that relative demand trends favoured higher educated workers during the 1990s, but reversed during the 2000s. We find that this trend reversal is key to understanding the observed patterns in earnings inequality in these economies. These results are robust to several alternative specifications, including taking into account the role of the minimum wage and cyclical conditions of the labor market.

JEL Codes: J20, J31
Keywords: Earnings Inequality, Unconditional Quantile Regressions, Supply-Demand Framework, Human Capital

---

*The views expressed in this article are those of the authors and not those of the Inter-American Development Bank.

†University of Oxford: manuel.fernandezsierra@economics.ox.ac.uk

‡Research Department. Inter-American Development Bank and IZA: julianm@IADB.ORG

1 Introduction

Latin America was the most unequal region in the world at the turn of the twentieth century (Alvaredo and Gasparini, 2015). However, in sharp contrast with previous trends, the first decade of the twenty first century was a period of rapid inequality decline for the region (Ferreira et al., 2008; Kahhat, 2010; López-Calva and Lustig, 2010; Gasparini and Lustig, 2011; Gasparini et al., 2011; Levy and Schady, 2013; Lustig et al., 2013). Redistribution through progressive fiscal policy, the emergence of conditional cash transfer programs to the poor, and changes in household demographics played a role in this transition. However, the most important contribution to inequality reduction was the decline in earnings inequality (Lopez-Calva and Lustig, 2010; Azevedo et al., 2013).

Flat or increasing wage inequality trends since the 1980s started to reverse in the late 1990s, albeit with varying turning points and intensity across countries. This contraction of the earnings distribution has been quite significant. For example, between 2000 and 2013, the 90th/10th interquartile range declined by 20 percent in Argentina, 28 percent in Chile, and a remarkable 46 percent in Brazil. Messina, Silva, and Lopez-Calva (2016) show that, since the start of the century, earnings inequality declined in 16 of the 17 countries in Latin America for which consistent statistics can be calculated.

In parallel with the change in inequality trends, most countries in the region also registered a rapid socio-demographic transformation in at least three dimensions: educational upgrading, age composition, and increased participation of women in the labor market. For example, the share of college educated workers since the early 1990s went from 16.5% to 26.6% in Argentina, almost doubled in Chile (14.3% to 27.7%), and almost tripled in Brazil (7.5% to 19.5%). Furthermore, the average worker age increased by more than a year in Argentina (37.3 to 39.0), by three years in Brazil (34.1 to 37.4), and by more than four years in Chile (35.8 to 40.4).

Our analysis starts by investigating how these changes in the skill-demographic composition of the workforce have shaped the wage structure in Argentina, Brazil and Chile during the last two decades. Following the work of Firpo et al. (2007) and Firpo et al. (2009) we construct counter-factual wage distributions to decompose the changes observed into changes in the composition of the labor force and changes in the pay structure.

We find that falling returns to education and labor market experience of the workforce explain a large share of the observed changes in inequality. Interestingly, the declining experience premium had a much larger explanatory power in the decline of inequality in the upper half of the distribution. In contrast, the decline of the returns to schooling explain a much larger share of inequality reduction at the bottom. Against these dominating patterns, pure composition changes related to the increase in educational attainment were inequality enhancing. This phenomenon, previously labeled as the paradox of progress

(Bourguignon et al., 2005), by which increases in educational attainment can be inequality increasing due to the convexity of the returns to education even when these changes are equalizing, appears to be still present in the region.

If different levels of human capital are imperfect substitutes in production, relative changes in educational attainment and labor market experience of the labor force, and relative changes in the returns to education and experience, should move in opposite direction. Hence, our analysis continues by assessing the role of labor supply changes in the observed education and experience premiums. Following the seminal work of Katz and Murphy (1992), Murphy and Welch (1992) and Card and Lemieux (2001), we build a stylized model of supply and demand for labor in which workers with different skill-demographic characteristics are imperfect substitutes in production. We then use household level data from Argentina, Chile and Brazil covering the last 25 years to estimate the parameters of the model and derive implications for the role of supply and demand factors in the evolution of experience and education premiums.

We show that a combination of imperfect substitutability across labor types and the observed movements in relative supplies go a long way at explaining the changes in the relative returns, especially the declines of both the schooling and experience premiums. Most of the fall in the high school/primary schooling premiums, which we find is a significant factor behind the fall in lower tail inequality in these countries, can be accounted for by the significant increase of workers with at least a high school degree. Moreover, according to our model the observed changes in labor supply should have resulted in an even greater reduction of the high school premium than the observed one, specially during the 1990s. This is because the demand for high-school workers increased in this period. Our results for the demand of college educated workers is similar. Rising supply of college educated workers has also pushed the college premium downwards over the past 25 years. We find that relative demand for college educated workers increased during the 1990s, but that this trend reversed since the start of the 2000s. The implication is that demand side trends attenuated the fall in the college premium during the 1990s, but accentuated the decline in the 2000s.

Further, we show that changes in the educational premiums are not the only factors driving the reconfiguration of the wage structure. The experience premiums also declined substantially, specially within groups of workers with similar levels of schooling, and this also contributed to inequality reduction. We provide novel estimates in the region for the elasticities of substitution between workers with different experience levels. These estimates suggest that aging of the workforce has contributed to changes in the experience premium, and through this channel to changes in the wage structure.

Our work is more closely related to that of Manacorda et al. (2010), which use a similar supply and demand framework to analyze changes in the wage structure in Latin

America during the 1990s. We depart from Manacorda et al. (2010) in two significant respects. First, our model allows for imperfect substitutability between workers with different potential experience within schooling levels. We show that this distinction is empirically relevant, and is one of the reasons they find workers of different age groups to be perfect substitutes in production, a feature rejected by the data in our model. Second, we extend the model to allow for differential demand trends within low skilled workers and experience groups. Our estimates suggest that relative demands favored more educated and experienced workers during the 1990s, but that there was a shift around the 2000s. Allowing for flexible demand trends becomes necessary to rationalize the patterns in the data. Beyond these two points, we reinforce their message of the importance of differentiating between workers with secondary schooling and those with primary schooling when thinking about labor market outcomes in Latin America, since we find strong evidence of imperfect substitutability between them.

The rest of the paper is organized as follows. Section 2 discusses the data and main stylized facts, reviewing the evolution of inequality and socio-demographic changes in Argentina, Brazil and Chile. Section 3 shows how changes in inequality were affected by compositional changes and changes in the wage structure associated with education, experience and gender. In section 4 we develop a simple stylized model of supply and demand based on the descriptive trends in the data; and Section 5 provides estimates to the key parameters of the model and a discusses their implications for the evolution of inequality. We provide robustness exercises that aim at understanding the sensitivity of our results to the modeling choices in Section 6. Finally, Section 7 concludes.

2 Data and Stylized Facts

The information for Argentina, Brazil and Chile is derived from household surveys provided by the Center for Distributive and Social Studies at Universidad de la Plata (CEDLAS) and The World Bank. The surveys are collected by the respective statistical agencies in each country, but they go through an homogenization procedure done by CEDLAS and The World Bank as part of the Socio-Economic Database for Latin America and the Caribbean (SEDLAC) project. All surveys include information about general characteristics of the workers (gender, age, education) and their jobs (type of contract, labor earnings, hours worked). With the exception of Argentina where information is restricted to urban areas,[1] all other surveys are nationally representative.

Each survey includes a question asking workers for the total monetary income from labor in a reference period. This is the variable that we use throughout the paper to capture labor earnings. The variable is divided by the self-reported total number of hours worked

[1] Urban areas account for almost 90% of the total population in the Argentina in 2013.

to obtain hourly earnings. The series are converted into real terms using the consumer price index of the respective countries.[2] We have restricted the sample to individuals between the ages of 16 and 65, and only use earnings of full time workers (individuals that self-reported working for more than 35 hours in the reference week[3]). For further details on the characteristics of the surveys and the construction of the variables see Appendix A.

Most countries in the Latin American region have experienced a sharp contraction of the earnings distribution over the past two decades. Figure 2.1 shows the evolution of three interquantile log wage ratios for the 3 countries in our sample. We use three different ratios to capture changes in wage dispersion. Changes in overall wage dispersion are summarized by the 90/10 log wage differential; inequality at the top of the distribution is traced by the 90/50 log wage difference; and we refer to lower tail inequality when we evaluate inequality below the median (50/10 log wage gap). The figure summarizes three of the main characteristics of earnings inequality in Latin America: first, the levels of earnings inequality are among the highest in the world. In 2013, the 90/10 log earnings ratio was close to 1.7 in the three countries, so that the earnings of a worker at the 90th percentile of the distribution is more than 5.5 times what a worker at 10th percentile gets. As a point of comparison, the OECD average of the 90/10 interdecile ratio in 2012 was 4.[4]

Second, income inequality has been declining during the past two decades,[5] reversing the increasing trend documented for the 1980s and the first years of the 1990s.[6] The reversal of the trend started at different years in our sample of countries, with peaks around 1996 in Chile and 2002 in Argentina. Inequality in Brazil has been declining for the full period of our analysis. Albeit this heterogeneity, the contraction of the earnings distribution is quite significant. Starting from the year in which we observed the peak in inequality in each country, the ratio between the 90th and 10th percentiles contracted by 31 percent in Argentina, 33 percent in Chile, and a remarkable 60 percent in Brazil.

Finally, contrasting with the large body of literature showing that most of the changes in the wage structure in high income countries have happened at the top of the income distribution,[7] changes in wage dispersion in our sample are observed both above

[2]Due to inconsistencies found in the official Consumer Price Index in Argentina (see Cavallo (2013)), we use the information from PriceStats (`http://www.statestreet.com/ideas/pricestats.html`) to deflate nominal wages in this country.

[3]The average share of workers that reported working less than 35 hours per week is 15 percent in Chile and Brazil, and close to 28 percent in Argentina. We present alternative estimates of our main results including part-time workers in the robustness section of the paper.

[4]See `http://stats.oecd.org/`.

[5]See Ferreira et al. (2008); Kahhat (2010); López-Calva and Lustig (2010); Gasparini and Lustig (2011); Gasparini et al. (2011); Levy and Schady (2013); Lustig et al. (2013)

[6]See Cragg and Epelbaum (1996); Londoño and Szekely (2000); Sánchez-Páramo and Schady (2003); Behrman et al. (2007); Cornia (2010); Manacorda et al. (2010) among others

[7]See Katz and Autor (1999); Autor et al. (2005, 2008); Lemieux et al. (2009); Acemoglu and Autor (2011) and the references therein.

and below the median. In the three countries we observed a drop in both the 90/50 and 50/10 log wage ratio since inequality peaked. Moreover, if we look at the full span of our data, only in Argentina we see a small increase in the 50/10 log wage ratio since the initial levels of 1995.

Figure 2.1: Interquantile Log Wage Ratio by Country

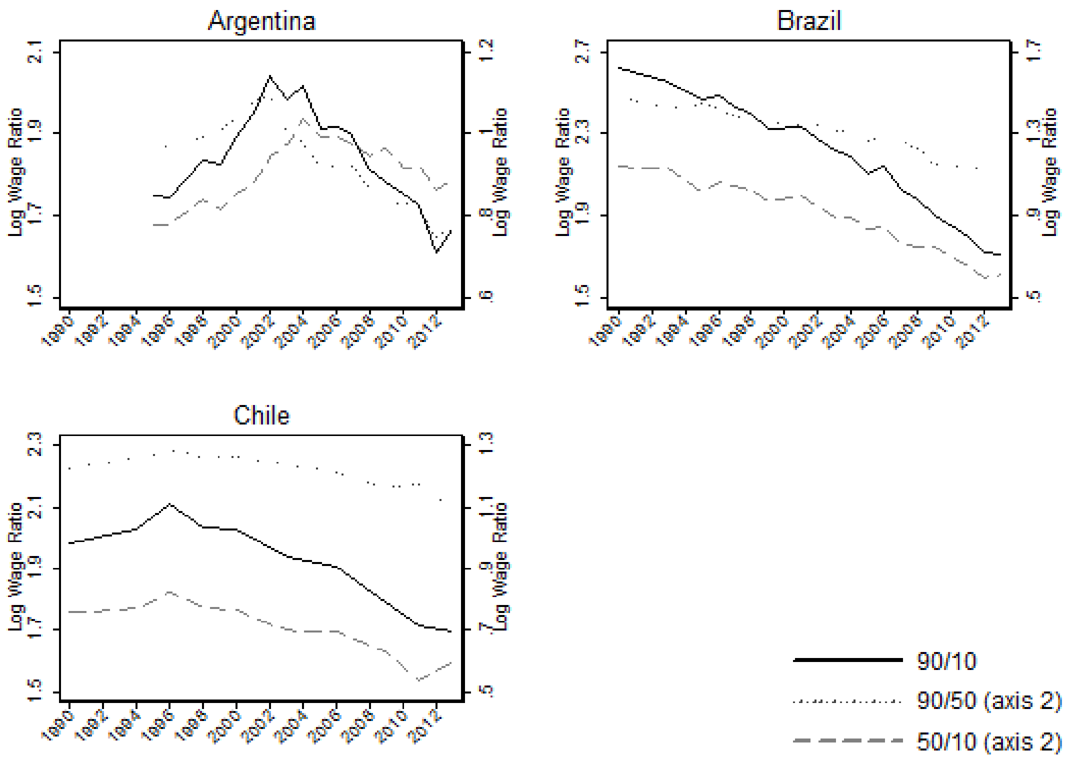

Notes: sample consists of full time workers (reported working 35 hours or more) between ages 16 and 65.

These changes in the distribution of earnings across the region are happening simultaneously to a considerable shift in the skill and demographic composition of the workforce. Table 2.1 shows the change in employment shares of different labor groups, defined based on their educational attainment, potential experience and sex. The table shows that there are at least three significant changes in the skill and demographic composition of employment during the last 20 to 25 years. First, there are sharp increases in the average level of education. The percentage of workers that reported having at most a primary education degree in the early 1990s was 47.2 percent in Argentina, 49.3 percent in Chile, and 77 percent in Brazil.[8] By 2013 these numbers had dropped substantially in each country, with the fall ranging between 17.7 percentage points in Argentina to 34.3 percentage points in Brazil. The gains in schooling are reflected in an increase of the share of workers with

[8] See Appendix A for details on the aggregation of workers with incomplete levels of schooling.

high school education completed, as well as by a rise in the share of workers with a college degree. For example, the share of college educated workers since the early 1990s went from 16.5% to 26.6% in Argentina; almost doubled in Chile (14.3% to 27.7%); and almost tripled in Brazil (7.5% to 19.5%).

The second major change in the composition of the workforce is related to the average age and experience levels of workers. The average age of a worker increased by more than a year in Argentina (37.3 to 39.0), three years in Brazil (34.1 to 37.4), and more than four years in Chile (35.8 to 40.4). Even with the sharp rise in the levels of schooling, this demographic shift has resulted in a rise in the average level of potential experience.[9] This is specially significant in the case of Chile, where the share of workers with more than 20 years of potential experience increased by 11.9 percentage points.

Finally, the third major change in the employment composition in these economies was the sharp rise in the participation of women in the labor market. By the early 1990s, female labor force participation was as low as 35 percent in Chile and close to 50 percent in Argentina and Brazil. By 2013 half of the women between the ages of 16 and 65 in Chile where working, and female labor force participation in Argentina and Brazil was close to 60 percent. As a consequence of this shift, the employment share of women in these economies rose substantially since the 1990s. In the cases of Brazil and Chile, the female employment share grew close to 8 percentage points. In the case of Argentina, where the initial participation was the largest among our sample of countries, the employment share of women increased by 2.3 percentage points.

[9] We define potential experience as: potential experience = age-years of education-6.

Table 2.1: 100 × Change in Employment Share

	Argentina	Brazil	Chile
	1995-2013	1990-2013	1990-2013
Education			
Primary or less	-17.71	-34.53	-24.17
High School	7.70	22.42	10.79
College	10.02	12.11	13.38
Pot. Exper.			
[≥20]	-0.22	2.01	11.93
Sex			
Female	2.32	8.88	8.26
Educ. + Pot. Exper.			
Primary or less [0-19]	-8.16	-20.60	-15.57
Primary or less ≥ 20	-9.55	-13.93	-8.60
High School [0-19]	3.37	12.24	-3.00
High School ≥ 20	4.33	10.18	13.79
College [0-19]	5.01	6.35	6.64
College ≥ 20	5.01	5.76	6.74

Notes: sample consists of full time workers (reported working 35 hours or more) between ages 16 and 65. Tabulated numbers are changes in the employment shares for each group.

3 Inequality, Workforce Compositional and Wage Structure

We move next to the question of the relative importance of changes in the composition of the labor force vs. changes in wage differentials across labor market groups to the recent evolution of inequality. A simple decomposition exercise can help disentangle the two. The idea is to exogenously fix the structure of relative wages at the average level across the last two decades, and quantify the counterfactual levels of the interquantile wage ratios under the observed compositional changes. Alternatively, we can keep the composition of the labor force fixed at a given point in time and construct counterfactual wage distributions to evaluate how changes in the schooling, experience and male premiums have affected the observed inequality dynamics.

The decomposition we propose follows Firpo et al. (2007, 2009), which have recently shown that using the properties of Recentered Influence Functions (RIF) one can extend the traditional Oaxaca-Blinder decomposition to analyze distributional statistics beyond the mean (e.g. quantiles). Details on the method are found in Appendix A.2. As a starting point, consider a transformed wage setting model of the form

$$RIFq_{\tau t} = X_t'\gamma_t + \epsilon_t \text{ for } t = 1, 0, \tag{3.1}$$

where t identifies the initial ($t = 0$) and final ($t = 1$) periods; $RIFq_{\tau t}$ represents the value of the RIF corresponding to the τ'th quantile of the earning distribution at time t; X is a vector of socio-demographic characteristics, which in our case contains a female dummy, years of education, years of education squared, years of potential experience, and years of potential experience squared. We can estimate equation (3.1) by OLS, and express the estimated difference over time of the expected value of the wage quantile $\hat{q}_\tau$ as

$$\Delta\hat{q}_\tau = \underbrace{\left(\overline{X'}_1 - \overline{X'}_0\right)\hat{\gamma}_P}_{\Delta\hat{q}_{X,\tau}} + \underbrace{\overline{X'}_P\left(\hat{\gamma}_1 - \hat{\gamma}_0\right)}_{\Delta\hat{q}_{S,\tau}}, \tag{3.2}$$

where overbars denote averages and $\hat{\gamma}_P$ and X_P correspond to the estimated vectors of parameters and the explanatory variables of a wage setting model in which observations are pooled across the two periods.[10] Here, $\hat{q}_{X,\tau}$ corresponds to the composition effect, which captures the part of the change in the τ'th wage quantile that is accounted by changes in the average skill-demographic characteristics of the workforce, given that we set the skill returns at their (weighted) average over the two periods; and $\hat{q}_{S,\tau}$ is the wage structure effect, and captures how changes in returns are affecting wages at the quantile τ, given that the observable characteristics are fixed to be equal to their (weighted) average over time.

Since we are interested in the effects of compositional and price changes on wage inequality, we construct the following measures for the 90/10, 90/50 and 50/10 log wage ratio in each country separately

$$\underbrace{\Delta\hat{q}_{90} - \Delta\hat{q}_{10}}_{\text{Overall}} = \underbrace{(\Delta\hat{q}_{X,90} - \Delta\hat{q}_{X,10})}_{\text{Composition}} + \underbrace{(\hat{q}_{S,90} - \hat{q}_{S,10})}_{\text{Wage Structure}} \tag{3.3}$$

$$\Delta\hat{q}_{90} - \Delta\hat{q}_{50} = (\Delta\hat{q}_{X,90} - \Delta\hat{q}_{X,50}) + (\hat{q}_{S,90} - \hat{q}_{S,50}) \tag{3.4}$$

[10] This specific counterfactual allows us to analyze composition and wage structure effects relative to a baseline defined by both the (weighted) mean returns and (weighted) mean characteristics over the two periods.

$$\Delta\hat{q}_{50} - \Delta\hat{q}_{10} = (\Delta\hat{q}_{X,50} - \Delta\hat{q}_{X,10}) + (\hat{q}_{S,50} - \hat{q}_{S,10}) \tag{3.5}$$

The results of the decomposition exercises are shown in Table 3.1. In the three countries we observe a very similar pattern: changes in the skill and demographic composition of the workforce have had an unequalizing effect on the distribution of earnings. Conditional on a counterfactual in which we fix the returns to observable characteristics at the (weighted) average level over the two periods, changes in the skill-demographic composition of the workforce would have resulted in a higher 90/10 log wage ratio. In Argentina this difference is small (5.6 log points) but it is sizable in Chile (24.9 log points), and Brazil (28.2 log points). These unequalizing effects of compositional changes are observed at both ends of the earnings distribution, but the magnitude tends to be larger at the upper tail (90/50 log wage ratio). An example of this is the case of Brazil and Chile, where, under our counterfactual scenario, the 90/50 log wage ratio would have increased in 22 and 18 log points, respectively.

A first clear message can be derived from the results of the decomposition exercises: changes in the skill-demographic composition alone cannot explain the observed patterns in income inequality in Latin America over the last two decades. Moreover, wage structure effects are having the dominant role. The alternative counterfactual exercise, where the skill-demographic composition of the workforce is fixed at the average level across periods, but we allow the price of those characteristics to change, predicts that the fall in inequality would have been even larger. In particular, wage inequality would have declined in an additional 33 percent had changes in the composition of the labor force be kept constant in Brazil and as much as an additional 60 percent in Argentina and 85 percent in Chile. Of course, these are partial equilibrium counterfactuals, which do not take into account the impact compositional changes may have had on the returns to observable characteristics, an aspect to which we come back later.

If we take a closer look at the wage structure effect for the three skill-demographic dimensions, we find that changes in the schooling premiums had a prominent role in the observed inequality trends in the three countries. The contribution of changes in the schooling premium outweighs the observed inequality decline in all three cases. Thus, changes in the schooling premium more than offset composition changes that were unequalizing during the period, and other unobservable factors that rowed against the declining inequality tide. This is particularly remarkable in Brazil and Chile, where, under our counterfactual scenario, changes in the schooling premium would have contributed to a decline in the 90/10 interquartile range of -115 and -168 log points (-68 and -81 percent), respectively.

But education was not the only aspect of human capital that contributed to the fall in inequality during the 2000s. Changes in the experience premium also played a significant role. In Argentina, the contribution of the decline in the experience premium (-28 log points) was as important as the contribution of the schooling premium (-27 log points). In Brazil and Chile the role of schooling was slightly larger, but the importance of the decline in the experience premium is also remarkable. Although changes in the gender wage gap had also equalizing effects, their impact on the overall inequality trends was much smaller.

Interestingly, lower and upper tail inequality appear to be driven by different factors. In particular, the change in the schooling premium is the fundamental factor behind the evolution of the 50/10 interquartile range, but it is only significant in the evolution of inequality in the upper half of the distribution in Chile. In contrast, changes in the experience premium are fundamental to understand upper half inequality. In Argentina and Brazil, they explain by themselves changes in the 90/50 interquartile range almost fully. In Chile, their role with respect to changes in the schooling premium is more modest, but are still a significant factor.

The decomposition exercise shows that the observed patterns in earnings inequality are mostly driven by how the wage structure is changing in time, but it gives no indication as to why those relative returns changed? Moreover, the wage structure effects are calculated under a counterfactual in which the skill-demographic composition of the workforce is held fixed, which we know was not the case. A natural hypothesis, then, is that the wage structure is changing because of the compositional changes, not despite of them. This would be the case if workers with different skill-demographic characteristics are not easily substituted in production, so that changes in relative supplies directly influence relative wages. In the next section we provide descriptive evidence that this simple mechanism is consistent with the observed trends, and proceed to formulate a model that could rationalize the patterns in the data.

Table 3.1: Compositional Changes and Inequality Patterns: Oaxaca-Blinder Decomposition Results

	Argentina (1995-2013)		Brazil (1990-2013)		Chile (1990-2013)	
	Est.	[S.E]	Est.	[S.E]	Est.	[S.E]
90/10						
Overall	-0.091	[0.016]	-0.850	[0.017]	-0.290	[0.021]
Composition	0.056	[0.007]	0.282	[0.016]	0.249	[0.011]
Education	0.054	[0.007]	0.302	[0.016]	0.211	[0.010]
Experience	0.001	[0.001]	-0.003	[0.001]	0.051	[0.002]
Sex	0.002	[0.001]	-0.017	[0.001]	-0.013	[0.001]
Wage Structure	-0.147	[0.018]	-1.132	[0.028]	-0.538	[0.023]
Education	-0.271	[0.113]	-1.153	[0.121]	-1.685	[0.086]
Experience	-0.282	[0.044]	-0.825	[0.095]	-0.497	[0.055]
Sex	-0.049	[0.009]	-0.042	[0.008]	-0.033	[0.007]
Constant	0.454	[0.139]	0.888	[0.230]	1.677	[0.120]
90/50						
Overall	-0.214	[0.014]	-0.350	[0.011]	-0.149	[0.018]
Composition	0.056	[0.005]	0.222	[0.008]	0.183	[0.009]
Education	0.056	[0.005]	0.229	[0.009]	0.169	[0.009]
Experience	0.001	[0.001]	-0.001	[0.001]	0.030	[0.002]
Sex	-0.001	[0.000]	-0.005	[0.001]	-0.015	[0.001]
Wage Structure	-0.270	[0.015]	-0.572	[0.015]	-0.332	[0.021]
Education	0.084	[0.059]	0.076	[0.041]	-1.021	[0.072]
Experience	-0.253	[0.036]	-0.204	[0.054]	-0.297	[0.053]
Sex	-0.020	[0.005]	-0.013	[0.004]	-0.019	[0.007]
Constant	-0.080	[0.079]	-0.431	[0.090]	1.004	[0.102]
50/10						
Overall	0.123	[0.018]	-0.500	[0.013]	-0.140	[0.017]
Composition	0.001	[0.005]	0.060	[0.013]	0.065	[0.006]
Education	-0.002	[0.004]	0.073	[0.014]	0.042	[0.006]
Experience	-0.000	[0.001]	-0.002	[0.001]	0.021	[0.001]
Sex	0.003	[0.001]	-0.011	[0.001]	0.002	[0.001]
Wage Structure	0.123	[0.018]	-0.559	[0.024]	-0.206	[0.019]
Education	-0.355	[0.101]	-1.228	[0.095]	-0.664	[0.067]
Experience	-0.029	[0.039]	-0.621	[0.052]	-0.201	[0.027]
Sex	-0.028	[0.008]	-0.029	[0.006]	-0.014	[0.006]
Constant	0.534	[0.124]	1.319	[0.166]	0.672	[0.085]

Notes: sample consists of full time workers (reported working 35 hours or more) between ages 16 and 65. Standard errors calculated via bootstrap with 100 replications.

4 The Role of Supply and Demand in the Evolution of Relative Returns

Figure 4.1 shows the evolution of the compositionally adjusted schooling premiums in each country for the three levels of education: primary, high school and college. The composition adjustment holds constant the relative employment shares of the different skill-demographic groups at their average levels across all years of the sample. In particular, we first compute mean (predicted) log real earnings in each country-year for 70 skill-demographic groups (5 education levels, 7 potential experience categories in five year intervals, males and females). Mean wages for broader groups shown in the figures are then calculated as fixed-weighted averages of the relevant sub-group means, where the weights are equal to the mean employment share of each sub-group across all years. This adjustment ensures that the estimated premiums are not mechanically affected by compositional shifts.

There are two main messages that can be derived from Figure 4.1. First, the trends in the compositionally adjusted schooling premiums follow very closely the patterns of the interquantile log wage ratios, specially for the high-school vs. primary wage gap. The peaks in overall (90/10) inequality coincide with the peaks in the high-school vs. primary wage gap in Argentina and Chile, and they are both falling since the start of the sample in Brazil. The magnitudes of these falls are also substantial. The relative premium of a worker with a high school degree vis-à-vis a worker with at most primary education declined (since the year in which overall inequality peaked) by 12 percent in Argentina, 19 percent in Chile, and 46 percent in Brazil. The college vs. high-school premiums are also falling during the same period, and in the case of Argentina this fall is even larger than the one observed for the high-school vs. primary wage gap (22% since 2002).

Second, the sharp fall in both of these schooling premiums is happening at the same time as the relative supply of more educated workers is increasing substantially. This is specially relevant when we focus on workers with the lowest schooling level. Figure 4.2 shows the average growth of the real log (compositionally adjusted) wages by education level, indexed so that each series takes the value of zero at the peak of inequality. Workers with at most primary education, which lost the most ground in terms of the employment share (see Table 2.1), gained substantially in terms of relative earnings. Albeit at a lesser extent, a similar pattern is observed when we compare workers with high school and college education, which, with the exception of Brazil, show a negative co-movement over time between their relative share in employment and the observed compositionally adjusted relative earnings.

The imperfect substitutability of high (college educated) and low (non-college educated) skill workers is a common assumption in the literature, but apart from the work of Manacorda et al. (2010),[11] there is little evidence of imperfect substitutability within workers at the lower education levels in Latin America. If this is the case, educational upgrading could be pushing up the relative earnings of primary educated workers by making them more scarce in the labor market.

Figure 4.1: Composition Adjusted College/High School and High School/Primary Wage Gap

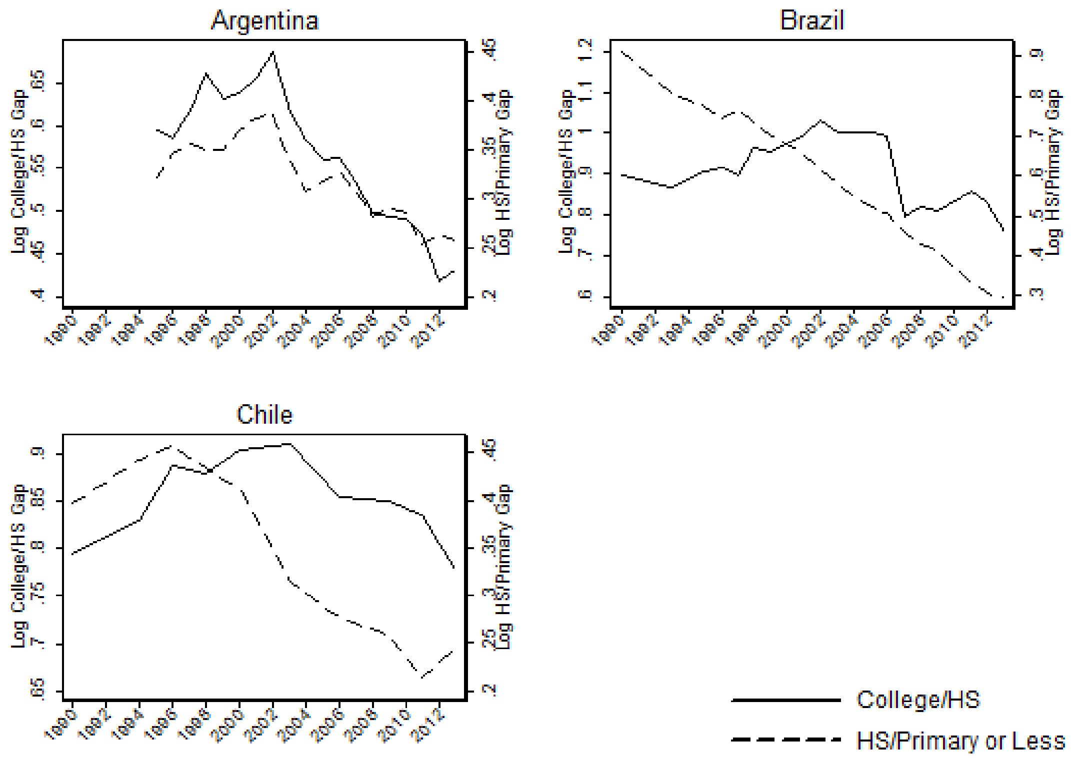

Notes: sample consists of full time workers (reported working 35 hours or more) between ages 16 and 65. See Appendix A for details on the construction of the compositionally adjusted series.

[11]The authors report an estimated elasticity of substitution between the two low skill groups in the order of 3.

Figure 4.2: Composition Adjusted Real Log Earnings by Education Level

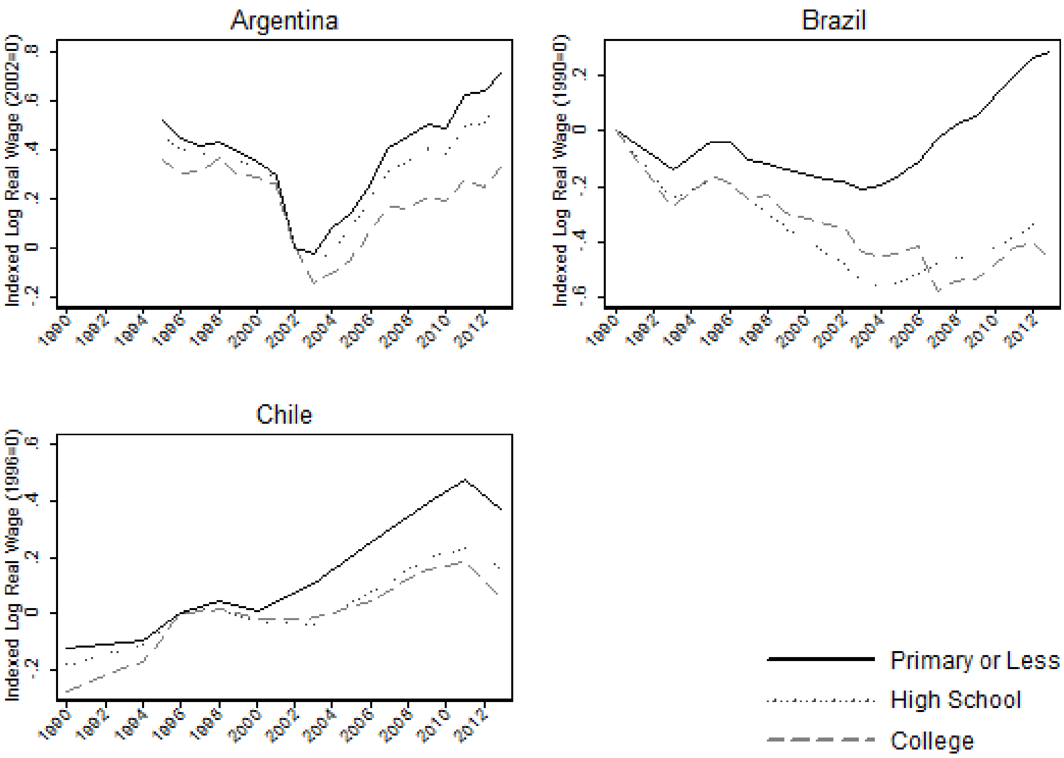

Notes: sample consists of full time workers (reported working 35 hours or more) between ages 16 and 65. See Appendix A for details on the construction of the compositionally adjusted series.

The negative co-movement between relative supplies and relative wages is less pronounced when we divide workers based on their potential experience, but compositional changes by experience level are also less dramatic than those observed by schooling. To see this, Table 4.1 shows the change over time of the compositionally adjusted log hourly earnings for different labor groups. In Chile, where the employment share of workers with more than 20 years of potential experience grew the most among the three countries (12 percentage points), the relative earnings vis-à-vis workers with less than 20 year of experience declined by 17 percent. Interestingly, the patterns appear to diverge conditional on education. The experience premium declined within each of the schooling levels in all countries, but relative supplies of more experienced workers only increased significantly within the low skill groups of Brazil and Chile, where the expansion of higher education is a more recent phenomenon (see Table 2.1).

Table 4.1: 100 × Changes in Real Composition Adjusted Log Hourly Earnings

	Argentina	Brazil	Chile
	1995-2013	1990-2013	1990-2013
All	12.23	0.58	37.58
Sex			
Male	12.38	-2.38	33.48
Female	11.87	6.33	45.69
Education			
Primary or less	20.09	28.94	48.83
High School	13.81	-32.71	33.76
College	-2.47	-46.76	32.09
Pot. Exper.			
[0-19]	15.81	-3.09	47.16
[≥20]	8.82	3.84	28.71
Educ. + Pot. Exper.			
Primary or less [0-19]	21.40	35.55	62.18
Primary or less ≥ 20	19.82	25.61	43.30
High School [0-19]	21.03	-24.02	47.94
High School ≥ 20	4.39	-49.48	16.21
College [0-19]	3.82	-41.01	37.53
College ≥ 20	-13.33	-55.00	22.68

Notes: sample consists of full time workers (reported working 35 hours or more) between ages 16 and 65. See Appendix A for details on the construction of the compositionally adjusted series.

Relative quantities and relative prices are moving in opposite directions for the two main drivers of the change in the wage structure: education and potential experience. How far can we go at explaining the trends in inequality by this simple mechanism? The answer to this question will depend on the sensitivity of the changes in relative wages to movements in relative supplies, that is, to the degree of substitutability between labor types with different skill-demographic characteristics. We now formalize this ideas with a stylized model of supply and demand in which a subset of the parameters of the model capture the key elasticities of substitution. We then proceed to estimate those parameters using the data from our three countries.

4.1 A Supply/Demand Model

The theoretical framework of the model follows the canonical work of Katz and Murphy (1992) and Murphy and Welch (1992).[12] We start by assuming that aggregate production in the economy can be described by a function that takes different labor inputs as its

[12] We deliberately borrow from Katz and Autor (1999) the way the theoretical framework is presented in this section.

arguments. Workers are divided into types, where each type is defined according to a set of skill-demographic characteristics. These types are allowed to be imperfect substitutes in production, so shifts in the supply or demand for workers with different characteristics can directly affect relative earnings. Factor prices and quantities are assumed to lay along the competitive equilibrium demand curve, so that each labor type is paid according to its marginal productivity. Under these conditions, labor input demands are a function of wages, the parameters of the production technology, and other "demand shifters",

$$L_t = D(W_t, Z_t), \tag{4.1}$$

where t indexes time periods; L_t is a $K \times 1$ vector of labor input demands; W_t is a $K \times 1$ vector of market wages; and Z_t is a $P \times 1$ vector of demand shifters. Possible sources of shifts in relative demand include non-neutral technical change, variations in non-labor input demands (e.g. through capital skill complementarity), product market demand shifts, trade, and outsourcing (Katz and Autor, 1999).

We can totally differentiate equation (4.1) to find the expression

$$dL_t = D_w dW_t + D_z dZ_t. \tag{4.2}$$

Assuming the production function is concave, the $K \times K$ cross-price demand matrix D_w is negative-semidefinite. This implies that changes in relative wages are negatively related to changes in relative supplies, once the effect of demand shifters is taken into account. That is:

$$dW_t = D_w^{-1}[dL_t - D_z dZ_t], \tag{4.3}$$

where,

$$dW_t' D_w dW_t = dW_t' \left[dL_t - D_z dZ_t\right] \leq 0. \tag{4.4}$$

This simple framework shows that in order to understand movements in relative wages across imperfectly substitutable groups of workers, we need to disentangle the effect of changes in relative supplies, which are observed in the data, from the effect of changes in relative demand, which are *not* observed in the data. Hence, to make the framework operational we need to add some structure to the model and specify both the general form of the production technology, and the way the demand shifters are allowed to enter the equations.

We assume that aggregate production in this economy can be described by a multilevel nested constant elasticity of substitution (CES) function. At the top level, output is produced as a CES combination of workers with high (college education completed or more) and low (high school degree at most) skills,

$$Y_t = \lambda_t \left(L_{Ut}^{\rho} + \alpha_t L_{St}^{\rho} \right)^{1/\rho}, \tag{4.5}$$

where Y_t is total output at time t; L_U is the total supply of low skill labor; L_S is the total supply of high skill labor; λ_t is a scale parameter that is allowed to vary in time to capture skill-neutral technological change; α_t is a time-varying parameter that captures both differences in relative productivities between skill and unskilled workers, and movements in relative demands between this two types; and ρ is a function of the elasticity of substitution (σ_ρ) between skilled and unskilled labor: $\sigma_\rho = \frac{1}{1-\rho}$.

As pointed out by Katz and Autor (1999), the fact that we model the economy using an aggregate production function means that we have to be careful not to interpret the parameters as if we where dealing with individual firms. For example, the elasticity of substitution σ_ρ reflects not only technical substitution possibilities between workers at the firm-level, but also outsourcing and substitution across goods and services in consumption. In a similar way, α_t captures relative productivity changes both at the intensive (workers performing better at the current jobs) and the extensive margins (e.g. a shift in work tasks across workers of different skill groups), changes in relative prices or quantities of non-labor inputs, and shifts in product demands among industries with different skill intensities.

At the second level of the production technology we further divide the total supply of unskilled labor (L_{Ut}) in two sub-groups. The first sub-group is formed by labor from workers that have at least obtained a high school degree, but that have not completed any post-secondary education. The second sub-group comprises labor from workers that have at most obtained a primary education degree. The aggregation is done using a productivity-weighted CES combination of the form

$$L_{Ut} = \left(L_{Pt}^{\delta} + \beta_t L_{Ht}^{\delta} \right)^{1/\delta}, \tag{4.6}$$

where L_{Pt} is the total supply of labor from workers with at most primary education; L_{Ht} is the total supply of labor from workers with at most secondary education; β_t is a time-varying parameter that captures both differences in relative productivities between the two sub-groups and changes in relative demands; and δ is a function of the elasticity of substitution (σ_δ) between the two low skill types.

Finally, we divide workers in each of the three schooling categories (primary, high school and college educated) in two potential experience sub-groups. The first sub-group is composed of workers that have less than 20 years of potential experience, henceforth denominated as inexperienced workers. The second sub-group comprises workers with 20 years of potential experience or more, henceforth denominated as experienced workers. In practice, we aggregate experience and inexperience workers within schooling levels using a productivity-weighted CES combination. In order to reduce the parameter space, we assume that the elasticities of substitution and the relative productivity parameters within the unskilled group (primary and high school educated) are the same. In particular, we have

$$L_{Kt} = \left(L_{KIt}^{\theta_U} + \phi_{Ut} L_{KEt}^{\theta_U}\right)^{1/\theta_U} \quad \text{for} \quad K = P, H \tag{4.7}$$

$$L_{St} = \left(L_{SIt}^{\theta_S} + \phi_{St} L_{SEt}^{\theta_S}\right)^{1/\theta_S}, \tag{4.8}$$

where I and E index inexperienced and experienced workers respectively; ϕ_{Ut} and ϕ_{St} are time-varying parameters that capture both differences in relative productivities and changes in relative demands between the potential experience sub-groups; and θ_U and θ_S are both functions of the elasticities of substitution (σ_{θ_U} and σ_{θ_S}) between the two experience sub-groups within the skilled and unskilled types.

There is a degree of arbitrariness in the way we specify the different nests of the production function, but our choices were made based on the observed patterns in the data. In particular, our production function allows both for imperfect substitutability between the two low skill sub-groups, and between potential experience types within workers of different schooling levels. To test the sensitivity of our estimates to this particular structure, we also estimate an alternative specification in the robustness section of the paper.

The demand side of the model has two types of relevant parameters that we wish to estimate: 4 parameters that are functions of the elasticities of substitution across types (ρ, δ, θ_U and θ_S), and a set of time varying relative productivities/demand shifters parameters (α_t, β_t, ϕ_{Ut} and ϕ_{St}). As shown by Johnson and Keane (2013), we could fit the trends in relative wages perfectly if we did not impose any restrictions on the evolution of the relative demand parameters, but this would mean that we would not be able to identify the parameters capturing the elasticities of substitution. We then restrict these relative productivities to follow a cubic trend in their natural logarithm.[13] For example, the

[13] We also tried both quadratic and quartic time trends, without significant changes to our main results

parameter α_t is allowed to change according to

$$\log \alpha_t = \alpha_0 + \alpha_1 \times t + \alpha_2 \times t^2 + \alpha_3 \times t^3. \tag{4.9}$$

We make two final assumptions that are key to our identification strategy. First, we assume that the labor supply of each labor type is exogenously determined. We acknowledge that this is a strong assumption, especially since we observe a sharp movement of women into the labor market during the last 20 years. Our choices for the relative wage and relative supply series discussed in the next section are made to ameliorate problems of selection arising from endogenous responses of women to changes in market conditions.

Second, we assume that the economy is operating along the competitive equilibrium demand curve. The implication of these assumptions is that the wage of each labor type is fully determined by it's marginal productivity. Given that we have 6 different labor types in the model (3 schooling levels $\times$ 2 potential experience groups), we get 6 equilibrium conditions. Denoting lower case variables as the natural logarithms of the respective upper case variables, the 4 equilibrium conditions for the low-skill types (PI, PE, HI and HE) are summarized in the following expression

$$\begin{aligned} w_{KJt} = \log \zeta_{KJt} &+ \frac{1}{\sigma_\rho}(y_t - l_{Ut}) + \frac{1}{\sigma_\delta}(l_{Ut} - l_{Kt}) \\ &+ \frac{1}{\sigma_{\theta_U}}(l_{Kt} - l_{KJt}) \quad \text{for} \quad K = H, P \quad \text{and} \quad J = E, I \end{aligned} \tag{4.10}$$

where $\zeta_{PIt} = 1$; $\zeta_{PEt} = \phi_{Ut}$; $\zeta_{HIt} = \beta_t$; and $\zeta_{HEt} = \beta_t \phi_{Ut}$. In a similar way, the two equilibrium conditions for the high-skill types (SI and SE) are

$$w_{SJt} = \log \zeta_{SJt} + \frac{1}{\sigma_\rho}(y_t - l_{St}) + \frac{1}{\sigma_{\theta_S}}(l_{St} - l_{SJt}) \quad \text{for} \quad J = E, I \tag{4.11}$$

where $\zeta_{SIt} = \alpha_t$; and $\zeta_{SEt} = \alpha_t \phi_{St}$.

(results are available upon request). The estimated parameters associated to the fourth order of the quartic specification were no longer statistically significant, while some of the third order coefficients were statistically significant. For this reason we chose cubic trends as our preferred specification. In the robustness section of the paper we also present the results from an excercise in which we allow for more flexibility in the specification of the time trends.

5 Results

5.1 Step I

The parameters of the model are estimated sequentially in three stages. In each stage we recover a subset of the parameters, and use them to construct the unobserved productivity weighted CES labor aggregates which are then used as inputs in the next step. For the first stage, we use the equilibrium conditions from equations (4.10) to find the expression that characterizes the evolution of relative earnings between experienced and inexperienced workers within the unskilled types. In particular, we have

$$w_{KEt} - w_{KIt} = \phi_{U0} + \phi_{U1} \times t + \phi_{U2} \times t^2 + \phi_{U3} \times t^3 - \frac{1}{\sigma_{\theta_U}}(l_{KEt} - l_{KIt}) \quad \text{for} \quad K = P, H. \tag{5.1}$$

Equation (5.1) says that changes in log relative wages between (low-skilled) experience groups depend on (1) the evolution of the log relative supplies, scaled by the inverse of the elasticity of substitution; and (2) the evolution of relative demand, as captured by the cubic time trends. Note that the relative earnings and relative labor supply series can be constructed directly from the data in each country. In all cases we limit our sample to population between ages 16 and 65. The labor supply of each labor type is equal to the total number of individuals with the respective socio-demographic characteristic across the entire working age population, irrespective of employment status or hours worked. Using working age population to construct the labor supply is more appropriate in our context than using the employed population, given the assumption that labor supply is exogenous in the model.[14]

The labor earnings series only uses full time workers (reported working 35 hours or more) when estimating the average earnings of the labor types.[15] Moreover, we further restrict the sample to include only male workers when constructing the relative earnings series. This is again done to address the concern of selection of women into the labor market.

We estimate the two equations in (5.1) by OLS pooling the data from the three countries. We allow the demand trends to be country-specific, but restrict the elasticities of substitutions to be common across countries. Both equations in (5.1) are estimated in a single regression that includes a skill dummy indicator (P/H). Results of the first step

[14]In the robustness section of the paper we show that using total employment or total hours worked has little effect on our estimates.

[15]We report the results when using both full-time and part-time workers in the robustness section of the paper.

estimates are shown in column 1 of Table 5.1. Our initial estimates provide evidence that workers of high and low experience within the unskilled group are not perfect substitutes, with the point estimate of the elasticity of substitution around 3.2.[16] Based on the observed changes in relative supplies in each country, the estimated elasticity of substitution implies a predicted fall in the experience premium, absent of any demand changes, of -6.6 percent in Argentina, -17 percent in Brazil, and -26 percent in Chile. This relative supply channel, by itself, closely matches the observed change in the experience premium within low skilled types in Chile (-25.6 percent), but underestimates the observed fall in Argentina (-12.3 percent), and slightly overestimates the observed decline in Brazil (-12.6 percent).

We further illustrate the negative co-movement of relative prices and relative quantities between experience levels in the Panel (a) of Figure 5.1. Here we show a scatter plot of log relative earnings and log relative supplies between experience types once the demand trends in equation (5.1) are accounted for. The log relative earnings series used in the Figure corresponds to the residuals of a regressions of the observed log relative earnings used in equation (5.1) on country specific cubic time trends and a skill dummy indicator. The log relative supply series used in the Figure correspond to the residuals of a regressions of observed log relative supplies on country specific cubic time trends and a skill dummy indicator. Once we depurate the series from changes in relative demands, the negative co-movement between relative earnings and relative supplies can be seen clearly.

We can also use the equilibrium conditions from equations (4.11) to arrive at a similar expression for the evolution of relative earnings between experienced and inexperienced workers within the high skilled types. This expression takes the form

$$w_{SEt} - w_{SIt} = \phi_{S0} + \phi_{S1} \times t + \phi_{S2} \times t^2 + \phi_{S3} \times t^3 - \frac{1}{\sigma_{\theta_S}}(l_{SEt} - l_{SIt}). \qquad (5.2)$$

The estimation of equation (5.2) is done in a similar way as it was done for the unskilled group, and results of the parameter estimates can be seen in column 2 of Table 5.1. We cannot reject the null hypothesis that experience and inexperience workers are perfect substitutes within the high skilled group, but the precision of the estimation is low, which can be partly explained by the small number of observations available for the regression. In the Panel (b) of Figure (5.1) we show a scatter plot of log relative earnings and log relative supplies between experience types once the demand trends in equation (5.2) are accounted for. The procedure used to remove the effect of changes in relative demands mimics the one described for the unskilled worker types.

[16] Using a different model specification for the United States, Card and Lemieux (2001) provide estimates of this elasticity between 4 and 6, while Johnson and Keane (2013) report estimates of around 10 which are not statistically significant.

Figure 5.1: Changes in Relative Earnings and Relative Supplies by Experience Level

(a) Unskilled Workers (b) Skilled Workers

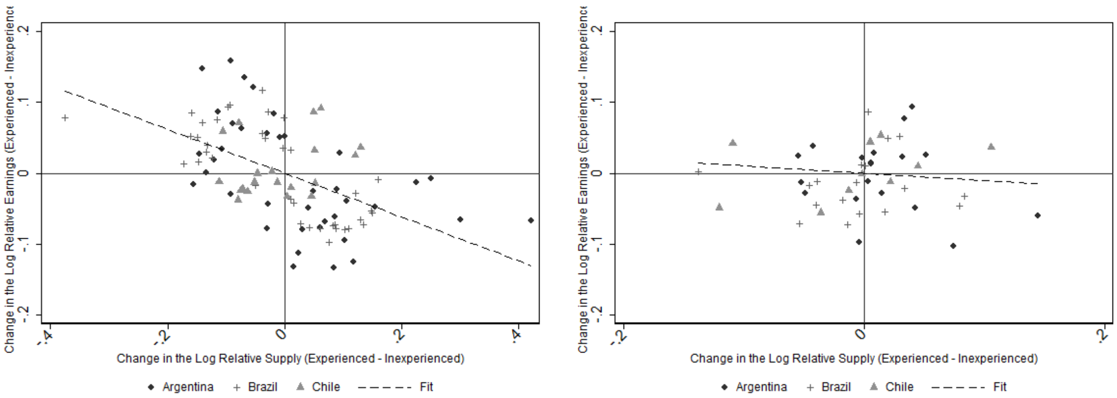

Notes: the Figure depicts changes in log relative earnings and log relative supplies once the country-specific demand trends in equation (5.1) and 5.2 are accounted for. Log relative earnings correspond to the residuals from a regression of observed relative earnings on country specific cubic time trends and a skill dummy. Log relative supplies correspond to the residuals from a regression of observed relative supplies on country specific cubic time trends and a skill dummy.

In this simple supply and demand framework, anything that is not explained by changes in relative supply ends up being explained by movements in relative demand, so this last component is identified residually. We show the evolution of the cubic demand trends captured by $\log \phi_{Ut}$ and $\log \phi_{St}$ in Figure 5.2.[17] The results are heterogeneous across countries. In Argentina and Chile, the results for both skilled and unskilled workers tend to show that relative demand for higher experience tended to increase during the 1990s, but has been either stagnant or in the decline since the beginning of the 2000s. The reversal of the trend is also observed in Brazil, but we only see a shift by the middle of the 2000s. The results suggests that the fall in the experience premiums where driven in part by the fact that the workforce in these economies are ageing, and that this mechanism was attenuated by a rise in relative demand for more experienced workers during the 1990s, but was later accentuated by a decline in relative demand during the 2000s.

[17]Each series is scaled so that it takes a value of zero at the first year in which data for the country is available.

Figure 5.2: Demand Trends Step I: Experienced vs. Inexperienced Workers by Skill Level

(a) Unskilled Workers ($\log \phi_{Ut}$) (b) Skilled Workers ($\log \phi_{St}$)

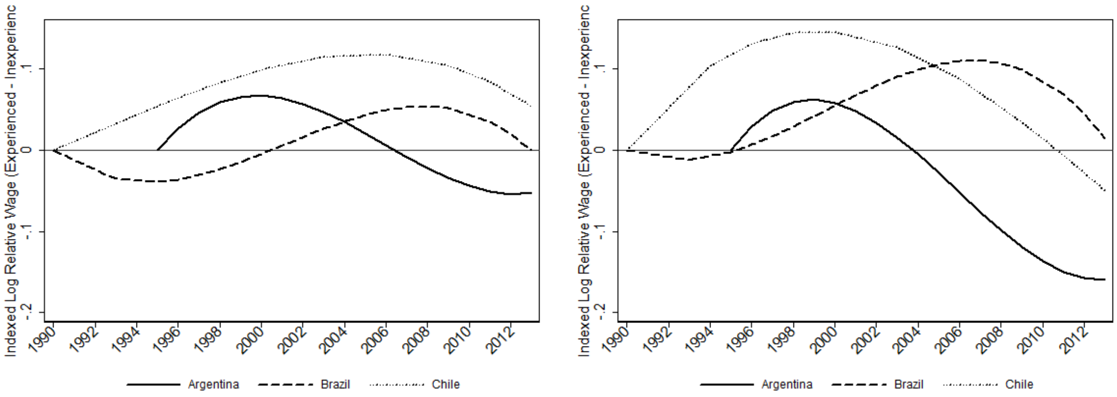

Notes: the Figure depicts the estimated relative demand trends between experience and inexperienced workers as captured by the country specific cubic time trends in equations (5.1) and (5.2). Each series is scaled so that it takes a value of zero at the first year in which data for the country is available.

5.2 Step II

The second step is aimed at recovering parameter estimates for the elasticity of substitution between the two low skill (non-college educated) types (σ_δ), and for the time trends capturing the evolution of their relative demands ($\log \beta_t$). We use the equilibrium conditions from equations (4.10) again, but this time we derive two equations characterizing the evolution of relative wages between workers with secondary education and those with at most primary education. In particular we have

$$\begin{aligned} w_{HJt} - w_{PJt} = \beta_0 + \beta_{1t} \times t + \beta_{2t} \times t^2 + \beta_{3t} \times t^3 - \frac{1}{\sigma_\delta}(l_{Ht} - l_{Pt})... \\ - \frac{1}{\sigma_{\theta_U}}\left[(l_{HJt} - l_{Ht}) - (l_{PJt} - l_{Pt})\right] \quad \text{for} \quad J = E, I \end{aligned}, \tag{5.3}$$

where both l_{Ht} and l_{Pt} are productivity-weighted CES labor aggregates. Although neither of the two labor aggregates is observed in the data, we can use the two equations in (4.7) and the estimated parameters from step I to calculate them. The first term after the cubic time trend in equation (5.3) is capturing overall (aggregated across experience groups) relative supplies between workers with at most primary education and workers with at most a high school degree. The second term after the cubic time trends represents relative changes in the potential experience composition between the two low skill groups. Note that the coefficient associated with this last term was already estimated in step I, so that we can use the results of this second step as a check for consistency.

We estimate both equations in (5.3) in a single regression, adding an experience group dummy indicator (E/I). As before, demand trends are allowed to be country-specific but the elasticities of substitution are assumed to be the same across countries. Results of the second step estimates are shown in column 3 of Table 5.1. Our estimated elasticity of substitution between workers with at most primary education and workers with at most secondary education is 2.2. This number is in line with the 2.8 point estimate found by Manacorda et al. (2010) for a different set of countries in the region during the 1990s, and reinforces the message that within the context of Latin America, there is a meaningful difference in the way the labor market treats the skills supplied by workers with secondary education and workers that only finish primary.

Moreover, most of the gains in schooling during the last 20 years in Argentina, Brazil and Chile have been the result of an increase in the share of the population with at least a high school education. This fact, combined with the imperfect substitutability between the lower skill sub-groups, implies that the compositional change in the schooling structure of the workforce are having large effects on relative earnings at the lower end of the distribution. To see this more clearly, Panel (a) of Figure 5.3 shows a scatter plot of log relative earnings and log relative supplies once the effect of relative demand trends and changes in the relative potential experience composition are taken into account. The log earning series is constructed as the residuals of an estimation of equation (5.3) that omits the aggregate relative supply term $(l_{Ht} - l_{Pt})$. The log relative supply series corresponds to the residuals of an estimation of equation (5.3) in which the aggregate relative supplies $(l_{Ht} - l_{Pt})$ are used as the dependent variable. The negative co-movement between relative supplies and relative earnings is clearly depicted in the Figure.

Panel (b) of Figure 5.3 shows the estimated demand trends as captured by $\log \beta_t$. The Figure shows that relative demand trends between workers in the lower schooling levels tended to favour the group with the higher educational attainment, specially in Argentina and Chile. A clear message from the results of the second step estimation is that the observed sharp declines in the high school/primary schooling premiums were mostly driven by the educational upgrading of the workforce, and that relative demand trends, if anything, tended to go against this pattern. This is clearly illustrated in Figure 5.4 that shows the fit of the model including or excluding demand trends. The exclusion of demand trends does not alter the model fit for Brazil, which is remarkably close to the observed relative wages. Instead, the decline of the high-school premium would have been larger in both Argentina and Chile had demand forces not favored high-school graduates against those with basic education.

Figure 5.3: Supply and Demand Factors Behind the Fall in the High-School/Primary Earnings Gap

(a) Changes in Log Relative Earnings and Log Relative Supplies

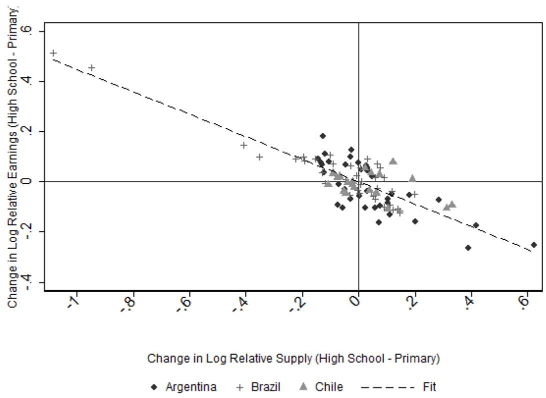

(b) Demand Trends Step II ($\log \beta_t$)

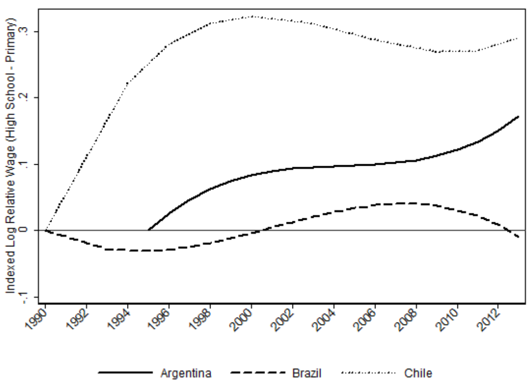

Notes: Panel (a) depicts changes in log relative earnings and log relative supplies between workers with at most a high school degree and those with only primary education, once the country-specific demand trends and changes in the relative potential experience composition are taken into account. The log earning series is constructed as the residuals of an estimation of equation (5.3) that omits the aggregate relative supply term $(l_{Ht} - l_{Pt})$. The log relative supply series corresponds to the residuals of an estimation of equation (5.3) in which the aggregate relative supplies $(l_{Ht} - l_{Pt})$ are used as the dependent variable. Panel (b) depicts the estimated relative demand trends between workers with at most a high school degree and those with only primary schooling as captured by the country specific cubic time trends in equation (5.3). Each series is scaled so that it takes a value of zero at the first year in which data for the country is available.

Figure 5.4: Goodness of Fit from the Second Step Estimates: High School/Primary Log Relative Earnings

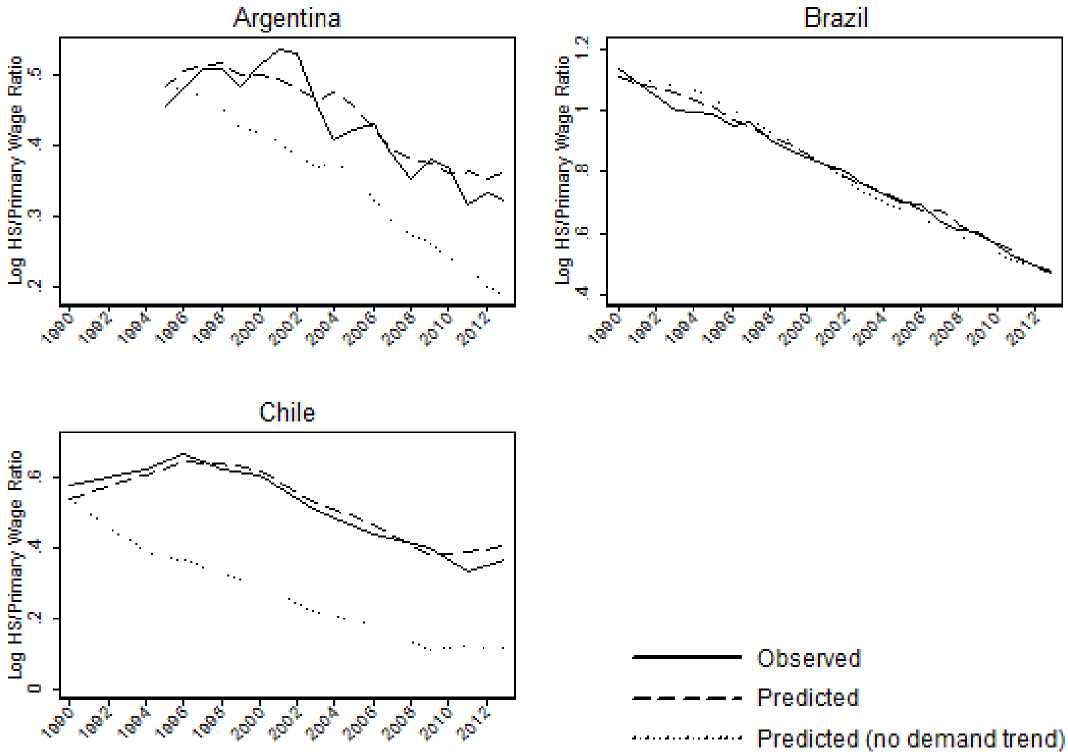

Notes: The Figure depicts the observed high school/primary schooling premium and the model prediction derived from the estimation of equation (5.3). An additional series constructed by setting the coefficients of the time trend to zero is included for comparison.

5.3 Step III

The goal of the last step is to obtain an estimate of the elasticity of substitution between skilled and unskilled workers (σ_ρ), and an estimate of the cubic trend describing their relative demand (α_t). With some manipulation of the equilibrium conditions from equations (4.10) and (4.11) we can derive the following 4 expressions,

$$\begin{aligned} \log(\frac{\widetilde{W_{SJt}}}{W_{KJt}}) = \log \alpha_t - \frac{1}{\sigma_\rho} \log(\frac{L_{St}}{L_{Ut}}) - \frac{1}{\sigma_\delta} \log(\frac{L_{Ut}}{L_{Kt}}) - \frac{1}{\sigma_{\theta_S}} \log(\frac{L_{SIt}}{L_{St}})... \\ - \frac{1}{\sigma_{\theta_U}} \log(\frac{L_{Kt}}{L_{KJt}}) \quad \text{for} \quad K = H, P \quad \text{and} \quad J = E, I \end{aligned} \qquad (5.4)$$

The left hand side of each of the equations in (5.4) consist of two terms: (1) the log relative earnings between skilled and unskilled workers of a given experience group; and

(2) different combinations of the relative demand parameters.[18] We put together those two terms using the symbol "$\sim$" to facilitate notation. In practice, the left hand side is simply a log relative earnings series that has been "demand-detrended" using the time trend estimates from the previous steps. The right hand side of the equations in (5.4) consists of the time varying parameter ($\log \alpha_t$) capturing movements in relative demand between skilled and unskilled workers; and a set of relative supply terms accompanied by their respective elasticity of substitution parameter. For example, the log relative supply between the CES aggregates of skilled and unskilled workers is multiplied by the inverse of the elasticity of substitution between those two groups. The same is true for the other three terms.

With the exception of L_{Ut} and L_{St}, all of the productivity-weighted CES labor aggregates have been previously estimated. Both L_{Ut} and L_{St} are defined by equations (4.6) and (4.8) respectively. We can use the parameter estimates from the previous steps to construct them, completing the set of variables needed for the estimation. Finally, note that each of the elasticities is identified using the variation in the relative supplies between the relevant labor groups.

The set of equations in (5.4) incorporate all the parameters of the production function. The estimation of these set of equations provides a third estimate for the elasticity of substitution σ_{θ_U}; a second estimate for both elasticities of substitution σ_{θ_S} and σ_{θ_ρ}; and a first estimate of the elasticity of substitution between skilled and unskilled workers σ_ρ. We estimate the four equations in a single regression, using a country-specific cubic demand trend for $\log \alpha_t$, and including skill-experience dummy indicators as covariates.

Results of the third step estimates are shown in column 4 of Table 5.1. The estimated elasticity of substitution between skilled and unskilled workers is 2.1, very close to the elasticity of substitution between the two low skill sub-groups. Our point estimate for this elasticity is higher than the 1.4, 1.6 and 1.64 values reported by Katz and Murphy (1992), Johnson and Keane (2013) and Autor et al. (2008) respectively for the United States; it is in line with the 2-2.5 range estimated by Card and Lemieux (2001) also in the United States; and is somewhat lower than the 2.5-5 range reported by Manacorda et al. (2010) for the Latin American region.

Once again, given the large rise in the share of college educated workers in these countries over the last 20 years, the results suggest that compositional changes in schooling are also having a strong impact on the college premium. To see this more clearly, Panel (a) of Figure 5.5 shows a scatter plot of log relative earnings and log relative supplies once the country-specific demand trends, changes in relative potential experience composition, and

[18] In particular, $\log(\widetilde{\frac{W_{SEt}}{W_{HEt}}}) = \log(\frac{W_{SEt}}{W_{HEt}}) - \log \frac{\hat{\phi}_{st}}{\hat{\beta}_t \hat{\phi}_{Ut}}$; $\log(\widetilde{\frac{W_{SEt}}{W_{PEt}}}) = \log(\frac{W_{SEt}}{W_{PEt}}) - \log \frac{\hat{\phi}_{St}}{\hat{\phi}_{Ut}}$; $\log(\widetilde{\frac{W_{SIt}}{W_{HIt}}}) = \log(\frac{W_{SIt}}{W_{HIt}}) - \log \frac{1}{\hat{\beta}_t}$; and $\log(\widetilde{\frac{W_{SIt}}{W_{PIt}}}) = \log(\frac{W_{SIt}}{W_{PIt}})$.

changes in schooling composition within the unskilled group are taken into account.[19] The negative co-movement between relative supplies and relative earnings is clearly depicted in the Figure.

Panel (b) of Figure 5.5 shows the estimated relative demand trends between skilled and unskilled workers, as captured by $\log \alpha_t$ in equation (5.4). In the three countries we observe a similar pattern, albeit with different magnitudes. Relative demand tended to favour college educated workers during the 1990s, but this trend started to reverse during the early 2000s. The break in the relative demand trend appears to happen around the year 2002 in each country.

The importance of demand trends is also illustrated in Figure 5.6, which shows the fit of the model with and without demand trends. The model performs relatively poorly when the demand channel is shut down. In all cases, supply channels alone over predict the fall of the college premium during most of the period. It is a strong demand for skills in the 1990s that slowly reverses in the 2000s that explain the inverted U shape of the skill premium. Thus, we conclude that while relative supply changes tended to push the college premium downwards over the past 20 years, relative demand changes ameliorated this effect during the 1990s, and magnified the relative supply channel from the new millennium onwards.

[19]The log earning series is constructed as the residuals of an estimation of equation (5.4) that omits the aggregate relative supply term ($\log(L_{St}/L_{Ut})$). The log relative supply series corresponds to the residuals of an estimation of equation (5.4) in which the aggregate relative supplies ($\log(L_{St}/L_{Ut})$) are used as the dependent variable.

Figure 5.5: Supply and Demand Factors Behind the Fall in the Skilled/Unskilled Earnings Gap

(a) Changes in Log Relative Earnings and Log Relative Supplies

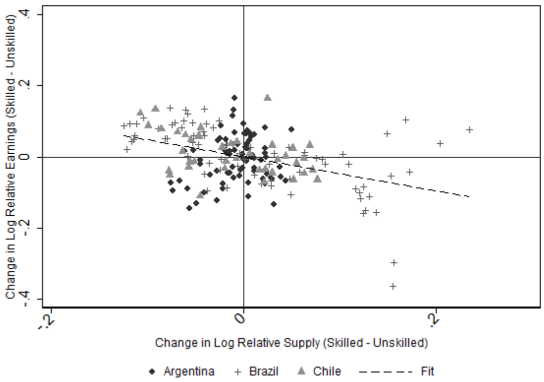

(b) Demand Trends Step III ($\log \alpha_t$)

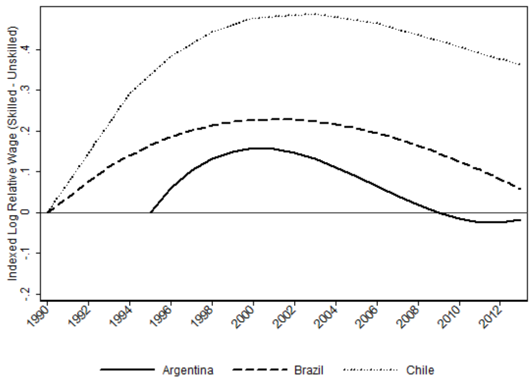

Notes: Panel (a) depicts changes in log relative earnings and log relative supplies between skilled and unskilled workers once the country-specific demand trends, changes in relative potential experience composition, and changes in scholling composition within the unskilled group are taken into account. The log earning series is constructed as the residuals of an estimation of equation (5.4) that omits the aggregate relative supply term $(l_{St} - l_{Ut})$. The log relative supply series corresponds to the residuals of an estimation of equation (5.4) in which the aggregate relative supplies $(l_{St} - l_{Ut})$ are used as the dependent variable. Panel (b) depicts the estimated relative demand trends between skilled and unskilled workers as captured by the country specific cubic time trend in equation (5.4). Each series is scaled so that it takes a value of zero at the first year in which data for the country is available.

Figure 5.6: Goodness of Fit from Third Step Estimates: Skilled/Unskilled Log Relative Earnings

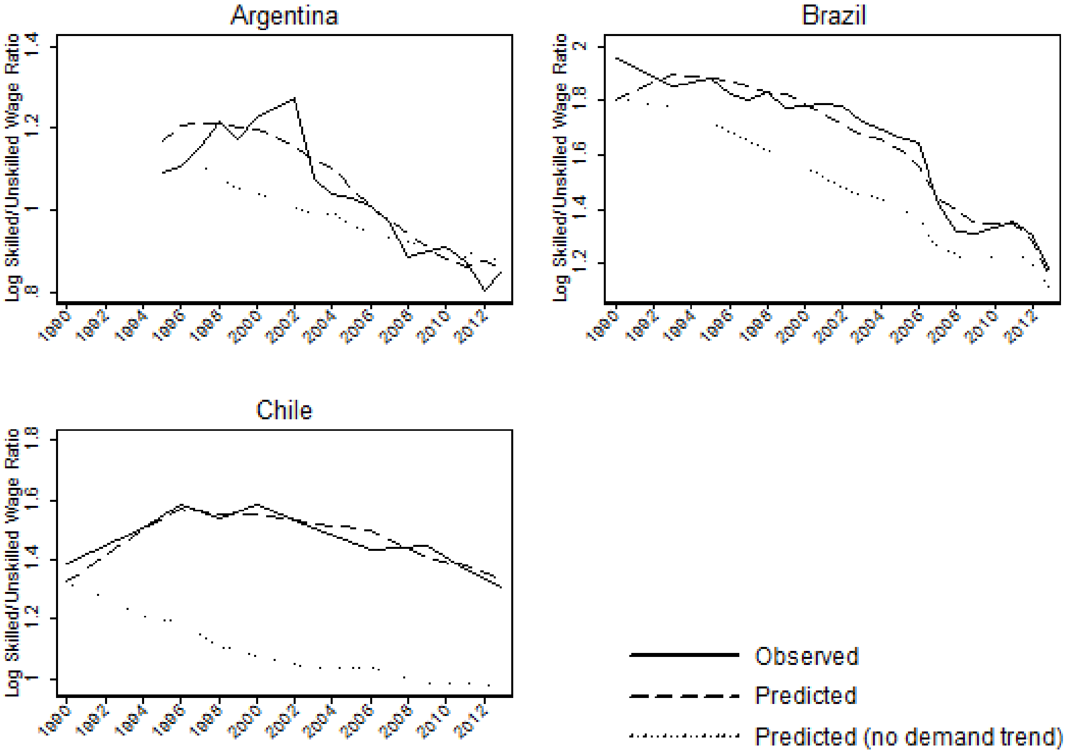

Notes: The Figure depicts the observed skilled/unskilled log earnings premium and the model prediction derived from the estimation of equation (5.4). The observed and predicted unskilled earnings series is constructed as a weighted average between the two low skill sub-groups, where the weights correspond to the respective labor share. An additional series constructed by setting the coefficients of the time trend to zero is included for comparison.

Table 5.1: Model Estimation Results

	STEP			
	IA	IB	II	III
Elasticities				
$-1/\sigma_{\theta_U}$	-0.309***		-0.323***	-0.359***
	(0.047)		(0.059)	(0.039)
$-1/\sigma_{\theta_S}$		-0.104		-0.111
		(0.191)		(0.081)
$-1/\sigma_{\delta}$			-0.448***	-0.471***
			(0.022)	(0.018)
$-1/\sigma_{\rho}$				-0.478***
				(0.125)
Demand Argentina				
Time	0.031**	0.034	0.029*	0.067***
	(0.015)	(0.025)	(0.016)	(0.014)
Time^2/100	-0.417*	-0.538*	-0.290	-0.830***
	(0.230)	(0.276)	(0.243)	(0.171)
Time^3/1000	0.127	0.165*	0.102	0.253***
	(0.096)	(0.094)	(0.100)	(0.061)
Demand Brazil				
Time	-0.019**	-0.011	-0.016*	0.045**
	(0.008)	(0.017)	(0.009)	(0.016)
Time^2/100	0.265**	0.246	0.218**	-0.251
	(0.082)	(0.216)	(0.096)	(0.178)
Time^3/1000	-0.079***	-0.086	-0.066**	0.029
	(0.023)	(0.069)	(0.026)	(0.051)
Demand Chile				
Time	0.011	0.035**	0.076***	0.094***
	(0.011)	(0.013)	(0.013)	(0.017)
Time^2/100	0.016	-0.234	-0.562***	-0.557***
	(0.146)	(0.174)	(0.157)	(0.154)
Time^3/1000	-0.022	0.032	0.124**	0.094**
	(0.045)	(0.053)	(0.048)	(0.041)
N	96	48	96	192
R^2	0.801	0.660	0.943	0.959

*** 1 percent ** 5 percent * 10 percent. Robust standard errors in parenthesis.

Notes: Each column presents the results of the estimation of the different stages of the model. Column IA shows the OLS estimates of the inverse of the elasticity of substitution between experience and inexperience workers within the low-skilled group ($\sigma_{\theta U}$) (see equation (5.1)); column IB correspond to the OLS estimates of the inverse of the elasticity of substitution between experience and inexperience workers within the skilled group ($\sigma_{\theta S}$) (see equation (5.2)); column II shows the OLS estimates of the inverse of the elasticity of substitution between the two low skill groups (σ_{δ}), and a second estimate of the inverse of the elasticity of substitution $\sigma_{\theta U}$ (see equation (5.3)); finally, column III shows the OLS estimates of the inverse of the elasticity of substitution between skilled and unskilled workers (σ_{ρ}), as well as additional estimates from the other elasticities in the model.

5.4 The Role of the Minimum Wage, Unemployment and the Commodity Price Boom

The simple supply and demand framework is silent about the role of institutional and cyclical conditions of the labor market in explaining the changes in the wage structure. This is a limitation that is specially relevant in our context since these economies have experienced sharp changes in labor market institutions like the minimum wage, and were exposed to very favourable external condition as a result of the commodity boom that started in the early 2000s (Erten and Ocampo, 2013; The World Bank, 2016). There is no straightforward way to incorporate these factors in this simple framework, but we follow Autor et al. (2008) and extend the empirical implementation of the model to examine the sensitivity of our estimates to including controls for changes in the minimum wage, the unemployment rate, and the commodity price boom.

Panel (a) of Figure 5.7 shows the evolution of the unemployment rate in Argentina, Brazil and Chile over the period of analysis. The three countries saw unemployment rising during the late 1990s and early 2000s, but these trends reversed since the year 2002. The pattern is specially clear in Argentina, which suffered a major economic and financial crisis between the end of 2001 and 2003. The cyclical conditions in these economies, as captured by the unemployment rates, coincide with the observe demand shifts that we have been documenting, especially for the college premiums. Panel (b) of Figure 5.7 shows the evolution of the real hourly minimum wage in the three countries, measured in USD of 2012. Between 1990 and 2012, the real minimum wage in Brazil and Chile increased by 120 percent and 138 percent respectively. In Argentina, the real minim wage remained flat between 1995 and 2002, but it increased in more than 130 percent over the ten year period that followed the economic meltdown. These numbers imply substantial real gains for workers at the lower end of the wage distribution, and are expected to be an alternative driving force behind the falls in the interquartile ratios.

The Panel (c) of Figure 5.7 shows the evolution of commodities prices since the early 1990s. South American countries are significant net commodity exporters. The commodity price super cycle that started in 2002 drastically improved the region's terms of trade (see Panel (d)), which paved the way for the expansion of domestic demand (The World Bank, 2016). The end result of the favourable external conditions was that South America's growth rate was around 5.5 percent on average between 2003 and 2013. The changes in relative labor demand at the start of the 2000s, documented in the previous section, could be a result of a favourable product demand shift in the commodity sector, which tends to be intensive in low skilled labor.

Figure 5.7: Labor Market Cyclical and Institutional Conditions

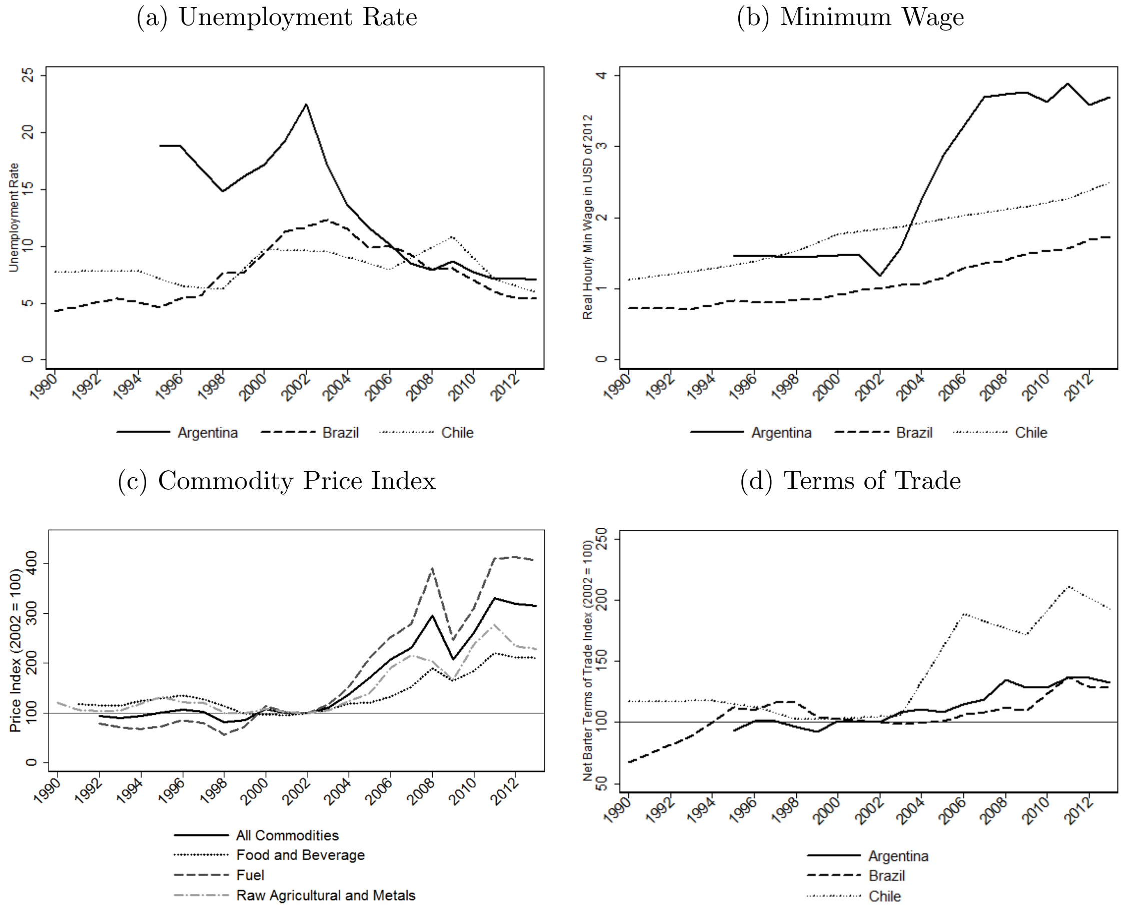

Notes: The unemployment rate series in Panel (a) is taken from the World Economic Outlook (WEO) database. The Minimum wage series in Panel (b) is taken from the annual indicators of the International Labour Organization (ILO). The source of the series in Panel (c) is The IMF's Primary Commodity Price System. The Food and Beverage series includes cereal, vegetable oils, meat, seafood, sugar, bananas, oranges, coffee, tea, and cocoa. The Fuel series includes crude oil (petroleum), natural gas, and coal. The Raw Agricultural and Metals series includes timber, cotton, wool, rubber, hides, copper, aluminum, iron ore, tin, nickel, zinc, lead, and uranium. See `http://www.imf.org/external/np/res/commod/index.aspx` for a description of the construction of the indices. The series in Panel (d) are taken from the United Nations Conference on Trade and Development (UNCTAD). Unit value indexes are based on data reported by countries, supplemented by UNCTAD's estimates using the previous year's trade values at the Standard International Trade Classification three-digit level as weights.

The extension of the empirical model is done by including the country specific natural logarithm of the minimum wage, unemployment rate, and the logarithm of the terms of trade index as covariates in our last step baseline result, corresponding to equation (5.4). To simultaneously identify changes in the demand trends, we restrict $\log \alpha_t$ to be common across countries, but allow for greater flexibility by estimating this relative demand using year and country fixed effects instead of the third order polynomial.

Column I of Table 5.2 shows the baseline model estimates when we allow for common relative demand trends. Results of this alternative specification are similar to our baseline model, but we observe a fall in the elasticity of substitution between skilled and unskilled workers, which goes from 2.1 to 1.5. Columns II-V show the results when we include the three additional set of covariates. The estimated coefficients associated to the minimum wage have a negative sign, with elasticities between -0.08 in Argentina to -0.50 in Chile for specification V. This numbers suggest that the rise in the minimum wage has contributed to the fall of the college premium, something that is expected. The fall of the unemployment rates during the past ten years tended to have an opposite effect, which is a puzzling result.

The improvements in the terms of trade appear to have contributed to the fall of the college premium, but the coefficients are not always statistically significant. The inclusion of these three variables, on the other hand, appears to add little in terms of the explanatory power of the model, and the relative supply channel is actually accentuated since once again we observe a fall in the elasticity of substitution between skilled and unskilled workers to about 1.28. Moreover, the exclusion of the relative supply variables leads to a significant fall in the R^2 as shown in Column VI.

These alternative specifications give some insights into what factors are behind the relative demand trend changes that we documented started in the early 2000s. Panel (a) of Figure 5.8 shows the evolution of the common demand trends as measured by the year specific fixed effects corresponding to Column I of Table 5.2. The trend reversal around the year 2002 in this baseline model is clearly depicted in the figure. In Panel (b) of Figure 5.8 we show the same demand trends for a model that includes controls for the log minimum wage and the unemployment rate. Once again, we observe a trend break during the same period of time, although the fall in the relative demand for high skill workers since 2002 is slightly attenuated.

A different story starts to emerge when we include controls for the change in the terms of trade in each country (see Panel (c)). In this case, we observe a deceleration of the relative demand for high skill workers around 2002, but it is significantly smaller than the one estimated in the baseline model. Moreover, if we include the full set of controls, the estimated relative demands remain flat after 2002. This results are suggestive that cyclical and institutional conditions of the labor market, specially regarding the improvements in external conditions of the economies following the commodity boom, where important determinants behind the acceleration of the fall of the skill premium during the 2000s. In the absence of this changes, demand trends favouring high skill workers would have attenuated the fall in earnings inequality brought about by the educational upgrading of the workforce.

Table 5.2: Model Estimation Results: Including and Excluding the Minimum Wage, the Unemployment Rate, and the Terms of Trade Index.

	Common Year Fixed Effects					
	I	II	III	IV	V	VI
Elasticities						
$-1/\sigma_{\theta_U}$	-0.272***	-0.290***	-0.287***	-0.283***	-0.301***	
	(0.038)	(0.037)	(0.037)	(0.037)	(0.034)	
$-1/\sigma_{\theta_S}$	-0.017	-0.102	-0.068	-0.060	-0.091*	
	(0.077)	(0.070)	(0.054)	(0.059)	(0.048)	
$-1/\sigma_{\delta}$	-0.446***	-0.442***	-0.448***	-0.446***	-0.448***	
	(0.017)	(0.014)	(0.014)	(0.014)	(0.012)	
$-1/\sigma_{\rho}$	-0.665***	-0.674***	-0.929***	-0.485***	-0.777***	
	(0.085)	(0.079)	(0.089)	(0.082)	(0.111)	
Log Real Min. Wage						
Argentina		-0.030			-0.083	-0.303
		(0.050)			(0.126)	(0.325)
Brazil		-0.189**			-0.380**	-0.591*
		(0.086)			(0.121)	(0.356)
Chile		0.064			-0.507**	-0.295
		(0.127)			(0.166)	(0.478)
Unemployment Rate						
Argentina			-0.009**		-0.008	-0.014
			(0.003)		(0.008)	(0.024)
Brazil			-0.025***		-0.027***	0.014
			(0.005)		(0.006)	(0.026)
Chile			0.006		0.016*	0.005
			(0.009)		(0.008)	(0.031)
Log Terms of Trade						
Argentina				-0.109	-0.592**	-0.373
				(0.109)	(0.239)	(0.681)
Brazil				-0.405***	-0.472***	0.230
				(0.089)	(0.102)	(0.757)
Chile				-0.007	-0.010	-0.167
				(0.048)	(0.121)	(0.340)
N	192	192	192	192	192	184
R2	0.967	0.972	0.973	0.971	0.978	0.791
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

*** 1 percent ** 5 percent * 10 percent. Robust standard errors in parenthesis.
Notes: the Table reports the third stage estimates of the parameters of the model when we include the natural logarithm of the real minimum hourly wage, the unemployment rate, and the natural logarithm of the terms of trade index. We include year and country fixed effects to capture relative demand trends.

Figure 5.8: Demand Trends Step III: including and Excluding Controls for Changes in the Minimum Wage, the Unemployment Rate and Commodity Prices.

(a) Year F.E.: Baseline

(b) Year F.E.: Controls for M.W. and U.R.

(c) Year F.E.: Controls for Terms of Trade

(d) Year F.E.: Full Set of Controls

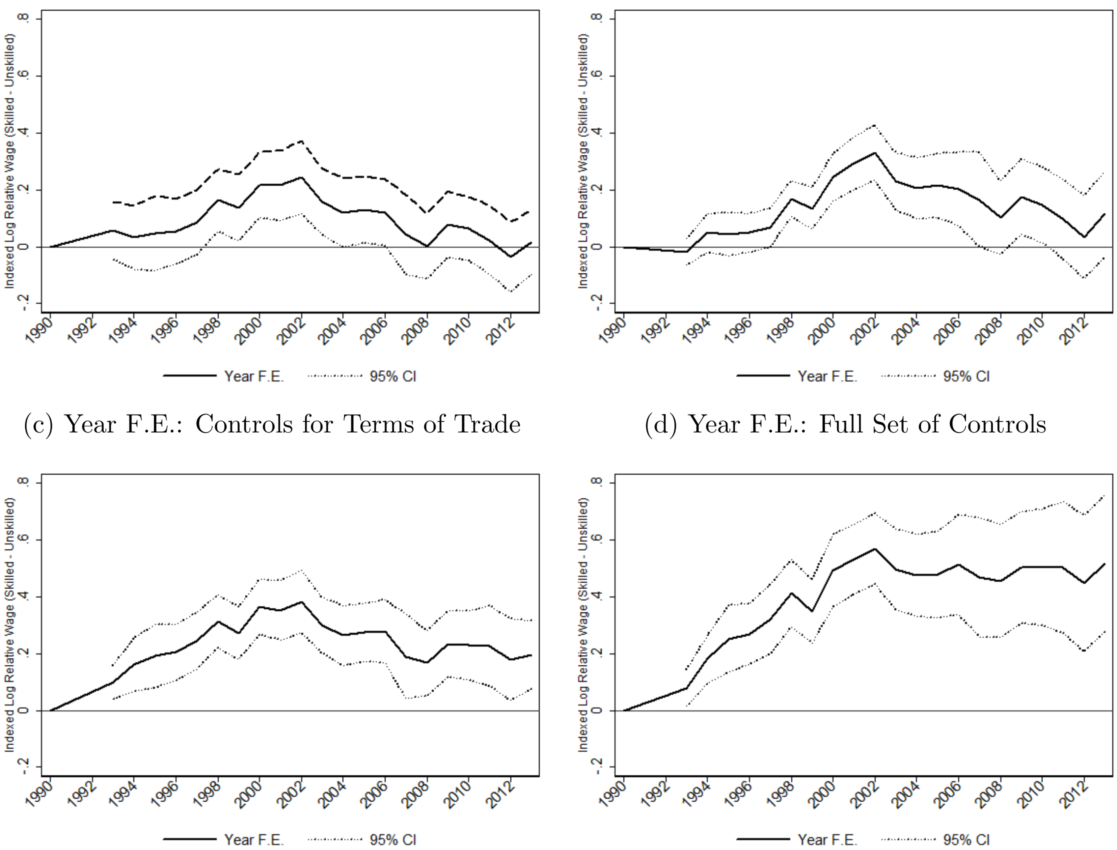

Notes: Each panel depicts the estimated relative demand trends between skilled and unskilled workers ($\log \alpha_t$) using different specifications of the last stage of the baseline model. Panel (a) corresponds to the estimates of the year fixed effects of column I in Table 5.2. The Panel (b) corresponds to the estimates of the year fixed effects of a model that includes controls for the unemployment rate and the natural logarithm of the minimum wage. Panel (c) corresponds to the estimates of the year fixed effects of column IV in Table 5.2, which includes controls for the log of the terms of trade index. Panel (d) corresponds to the estimates of the year fixed effects of column V in Table 5.2, which includes controls for the log real minimum wage, the unemployment rate of each country, and the log of the terms of trade index. The demand trends are scaled so that they takes a value of zero in 1990.

6 Robustness

Table 6.1 presents the last stage results of the model estimates using alternative supply measures. In our baseline specification (re-estimated in Column I), labor supply is calculated by adding up the total number of individuals of a given skill-demographic type between the ages of 16 and 65. We use this measure as our preferred specification because of our assumption that the supply of the skill-demographic groups is exogenous. In the second column of Table 6.1 we presents the results when we restrict the sample to include only occupied population in the supply measures. There are no significant changes in our estimates. In column 3 we present the same estimates when labor supply is calculated by adding up the total number of hours worked by each labor type. This is done to account for potential changes in the intensive margin. Since we only have information on hours worked for individuals that are occupied, we impute those numbers for individual outside the workforce by assigning them the average number of hours worked by an employed worker with the same education, potential experience, and sex in the respective country-year. The elasticity of substitution between skilled and unskilled workers increases from 2.1 in our main specification to 2.6 when we use hours worked, but the rest of the parameters' point estimates are unchanged.

Table 6.1: Model Estimation Results: Alternative Supply Measures

	Supply Measure		
	Working Age Pop.	Occupied Pop.	Total Hours Worked
Elasticities			
$-1/\sigma_{\theta_U}$	-0.359***	-0.337***	-0.321***
	(0.039)	(0.040)	(0.040)
$-1/\sigma_{\theta_S}$	-0.111	-0.128	-0.092
	(0.081)	(0.078)	(0.075)
$-1/\sigma_{\delta}$	-0.471***	-0.457***	-0.451***
	(0.018)	(0.018)	(0.018)
$-1/\sigma_{\rho}$	-0.478***	-0.505***	-0.393***
	(0.125)	(0.115)	(0.105)
Demand Argentina			
Time	0.067***	0.076***	0.075***
	(0.014)	(0.014)	(0.015)
$\text{Time}^2/100$	-0.830***	-0.958***	-0.920***
	(0.171)	(0.170)	(0.175)
$\text{Time}^3/1000$	0.253***	0.290***	0.274***
	(0.061)	(0.060)	(0.061)
Demand Brazil			
Time	0.045**	0.048**	0.055**
	(0.016)	(0.016)	(0.017)
$\text{Time}^2/100$	-0.251	-0.280	-0.289
	(0.178)	(0.171)	(0.178)
$\text{Time}^3/1000$	0.029	0.035	0.031
	(0.051)	(0.049)	(0.050)
Demand Chile			
Time	0.094***	0.088***	0.102***
	(0.017)	(0.015)	(0.015)
$\text{Time}^2/100$	-0.557***	-0.559***	-0.596***
	(0.154)	(0.139)	(0.143)
$\text{Time}^3/1000$	0.094**	0.105**	0.101**
	(0.041)	(0.038)	(0.039)
N	192	192	192
R^2	0.959	0.959	0.957

*** 1 percent ** 5 percent * 10 percent. Robust standard errors in parenthesis.

Notes: Each column presents the results of the estimation of the third stage of the model using alternative measures to construct the total labor supply. The first column corresponds to our baseline results; the second column limits the sample to include only occupied population; and the final column uses total hours worked.

Table 6.2: Model Estimation Results. Prime Age Workers (25-55)

	STEP			
	IA	IB	II	III
Elasticities				
$-1/\sigma_{\theta_U}$	-0.345***		-0.398***	-0.426***
	(0.049)		(0.042)	(0.044)
$-1/\sigma_{\theta_S}$		-0.182		-0.195**
		(0.148)		(0.068)
$-1/\sigma_{\delta}$			-0.439***	-0.462***
			(0.023)	(0.019)
$-1/\sigma_{\rho}$				-0.562***
				(0.101)
Demand Argentina				
Time	0.038**	0.055**	0.027	0.069***
	(0.018)	(0.022)	(0.018)	(0.016)
Time^2/100	-0.595**	-0.815**	-0.346	-0.916***
	(0.265)	(0.288)	(0.262)	(0.191)
Time^3/1000	0.206*	0.265**	0.140	0.293***
	(0.109)	(0.101)	(0.109)	(0.067)
Demand Brazil				
Time	-0.027**	0.009	-0.014	0.042**
	(0.010)	(0.015)	(0.009)	(0.017)
Time^2/100	0.351***	0.087	0.194*	-0.233
	(0.100)	(0.199)	(0.104)	(0.178)
Time^3/1000	-0.098***	-0.057	-0.055*	0.031
	(0.028)	(0.066)	(0.030)	(0.048)
Demand Chile				
Time	0.007	0.044**	0.081***	0.099***
	(0.012)	(0.015)	(0.013)	(0.014)
Time^2/100	0.030	-0.392**	-0.614***	-0.661***
	(0.151)	(0.178)	(0.161)	(0.126)
Time^3/1000	-0.020	0.082	0.139**	0.130***
	(0.046)	(0.052)	(0.049)	(0.035)
N	96	48	96	192
R^2	0.801	0.642	0.936	0.960

*** 1 percent ** 5 percent * 10 percent. Robust standard errors in parenthesis.

Notes: Each column presents the results of the estimation of the different stages of the model. Column IA shows the OLS estimates of the inverse of the elasticity of substitution between experience and inexperience workers within the low-skilled group ($\sigma_{\theta U}$) (see equation (5.1)); column IB correspond to the OLS estimates of the inverse of the elasticity of substitution between experience and inexperience workers within the skilled group ($\sigma_{\theta S}$) (see equation (5.2)); column II shows the OLS estimates of the inverse of the elasticity of substitution between the two low skill groups (σ_{δ}), and a second estimate of the inverse of the elasticity of substitution $\sigma_{\theta U}$ (see equation (5.3)); finally, column III shows the OLS estimates of the inverse of the elasticity of substitution between skilled and unskilled workers (σ_{ρ}), as well as additional estimates from the other elasticities in the model.

Figure 6.1: Demand Trends. Prime Age Workers (25-55)

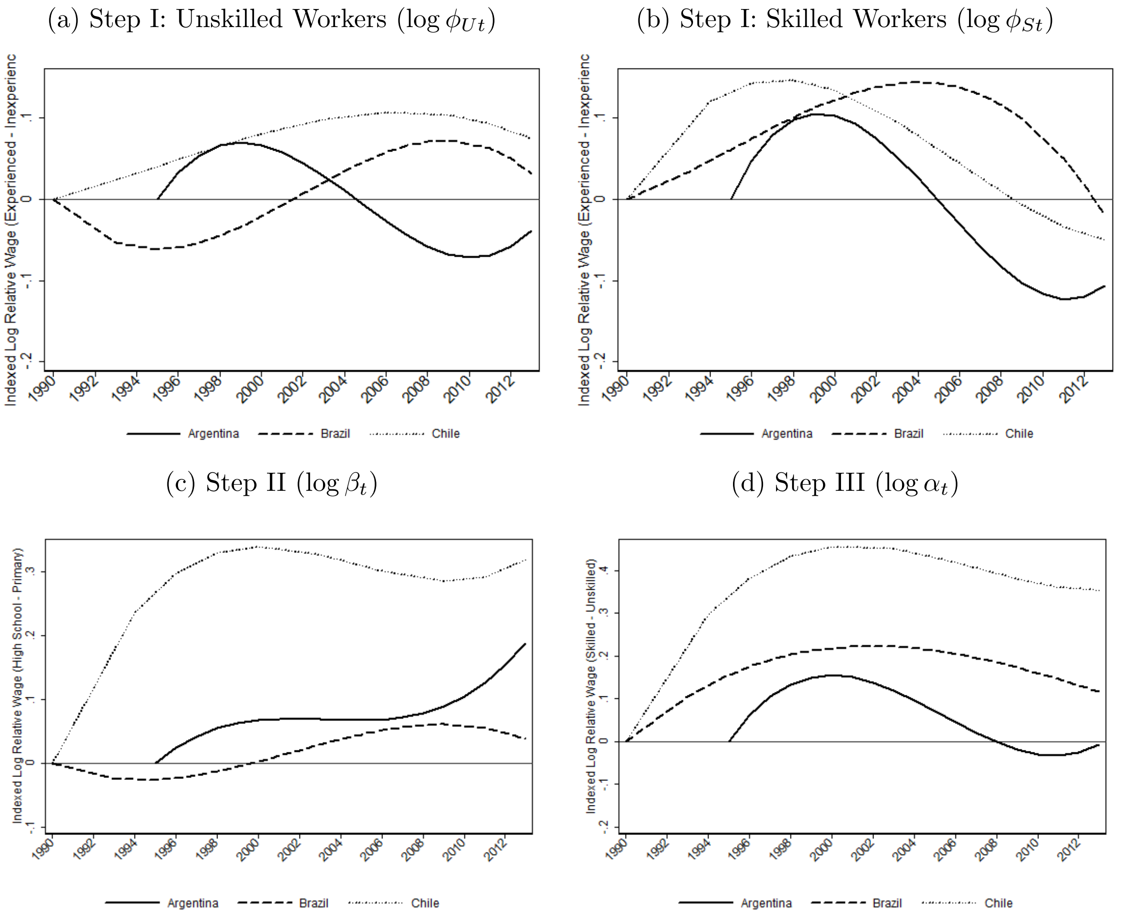

Notes: Panel (a) and (b) depict the estimated relative demand trends between experience and inexperienced workers as captured by the country specific cubic time trends in equations (5.1) and (5.2). Panel (c) depicts the estimated relative demand trends between workers with at most a high school degree and those with only primary schooling as captured by the country specific cubic time trends in equation (5.3). Panel (d) depicts the estimated relative demand trends between skilled and unskilled workers as captured by the country specific cubic time trend in equation (5.4). Each series is scaled so that it takes a value of zero at the first year in which data for the country is available.

Table 6.3 presents the third stage results of a second exercise in which we include both part-time and full-time workers in the calculation of the wage series. Once again, there is a small increase in the point estimate of the elasticity of substitution between high and low skilled workers when we use hours worked (2.7), but beyond this there are no major changes from our baseline results.

Table 6.3: Model Estimation Results: Part-Time and Full-Time Workers.

	Supply Measure		
	Working Age Pop.	Occupied Pop.	Total Hours Worked
Elasticities			
$-1/\sigma_{\theta_U}$	-0.389***	-0.373***	-0.359***
	(0.044)	(0.043)	(0.043)
$-1/\sigma_{\theta_S}$	-0.171*	-0.171*	-0.125
	(0.094)	(0.092)	(0.090)
$-1/\sigma_{\delta}$	-0.478***	-0.466***	-0.462***
	(0.020)	(0.020)	(0.019)
$-1/\sigma_{\rho}$	-0.432**	-0.467***	-0.372**
	(0.136)	(0.130)	(0.112)
Demand Argentina			
Time	0.040**	0.049**	0.049**
	(0.015)	(0.015)	(0.016)
$\text{Time}^2/100$	-0.414**	-0.544**	-0.501**
	(0.206)	(0.204)	(0.213)
$\text{Time}^3/1000$	0.125	0.165**	0.145*
	(0.080)	(0.078)	(0.080)
$\text{Time}^3/10000$			
Demand Brazil			
Time	0.049**	0.051**	0.058**
	(0.017)	(0.017)	(0.018)
$\text{Time}^2/100$	-0.301	-0.331*	-0.338*
	(0.197)	(0.192)	(0.193)
$\text{Time}^3/1000$	0.049	0.055	0.052
	(0.057)	(0.055)	(0.054)
$\text{Time}^3/10000$			
Demand Chile			
Time	0.103***	0.096***	0.110***
	(0.021)	(0.019)	(0.017)
$\text{Time}^2/100$	-0.702***	-0.691***	-0.736***
	(0.184)	(0.167)	(0.168)
$\text{Time}^3/1000$	0.150**	0.156***	0.155**
	(0.050)	(0.046)	(0.047)
$\text{Time}^3/10000$			
Observations	192	192	192
R^2	0.942	0.943	0.943

*** 1 percent ** 5 percent * 10 percent. Robust standard errors in parenthesis.
Notes: the Table reports the third stage estimates of the parameters of the model when we include part-time workers in the construction of the wage series. Each column presents the results using alternative measures of the total labor supply by each group. The first column corresponds to our baseline results; the second column limits the sample to include only occupied population; and the final column uses total hours worked.

The particular structure we use when setting up the demand side of the model is based on the observed patterns in the data. But there is a degree of arbitrariness in the modelling decisions, so that it is important to understand how sensitive are our results to alternative specifications. In the Appendix A.3 we present a different formulation of the demand side that follows the work of Manacorda et al. (2010), with some modifications to allow for comparability to our estimates. There are two main differences with respect to our baseline model that are worth pointing out: first, the ordering of second and third levels of the production function is reversed. In particular, skilled and unskilled workers are first disaggregated between 7 experience groups, and then further divided between the two lowest schooling levels within each potential experience category. This change implies that the elasticity of substitution between potential experience groups is the same for skilled and unskilled workers by construction.

Second, the number of potential experience groups is larger (7 in five year intervals), but the identifying assumption is that the relative productivities/demand shifters between experience sub-groups is time-invariant. This is an important difference given that our baseline results show that relative demand for workers with higher experience levels has not remained constant.

The parameter estimates of this alternative model are presented in Table 6.4.[20] The elasticity of substitution between the two low skill groups is larger than the 2.2 value estimated in our baseline result, with a magnitude that fluctuates between 2.68 and 3.78. These numbers are similar to the ones reported by Manacorda et al. (2010). Although not directly comparable, the point estimates in our model for the two elasticities of substitution between workers with different levels of potential experience ranged between 3 and 9, depending on the skill group. The analogous (single) estimate in the alternative specification is between 5.8 and 7.24. This new estimates are still statistically significant, but they imply that the sensitivity of relative wages to changes in relative supplies across experience groups is smaller in this set up, at least relative to the the low-skilled types. Finally, the elasticity of substitution between skilled and unskilled workers falls in this alternative model, going from 2.1 to 1.31. This number is closer to similar elasticities estimated in the United States (see, Katz and Murphy (1992) and Johnson and Keane (2013)), and would imply that the sensitivity of relative earnings to changes in relative supplies is even higher than what we found.

[20]See Manacorda et al. (2010) for a description of the estimation procedure.

Table 6.4: Model Estimation Results: Alternative Production Function

	STEP		
	I	II	III
Elasticities			
$-1/\sigma_\delta$	-0.372***	-0.264***	-0.341***
	(0.026)	(0.011)	(0.014)
$-1/\sigma_\theta$		-0.138**	-0.171***
		(0.045)	(0.043)
$-1/\sigma_\rho$			-0.762***
			(0.047)
Demand Argentina			
Time	0.027*		0.108***
	(0.014)		(0.012)
$\text{Time}^2/100$	-0.002		-0.011***
	(0.002)		(0.002)
$\text{Time}^3/1000$	0.000		0.000***
	(0.000)		(0.000)
Demand Brazil			
Time	-0.003		0.007
	(0.012)		(0.010)
$\text{Time}^2/100$	0.001		0.001
	(0.001)		(0.001)
$\text{Time}^3/1000$	-0.000		-0.000
	(0.000)		(0.000)
Demand Chile			
Time	0.062***		0.084***
	(0.010)		(0.009)
$\text{Time}^2/100$	-0.005***		-0.005***
	(0.001)		(0.001)
$\text{Time}^3/1000$	0.000**		0.000**
	(0.000)		(0.000)
N	336	672	672
R^2	0.973	0.934	0.886

*** 1 percent ** 5 percent * 10 percent. Robust standard errors in parenthesis.
Notes: Each column presents the results of the estimation of the different steps in the alternative model described in the Appendix A.3.

The estimated time trends captured by the relative demand parameters in the alternative model are presented in Figure 6.2. The observed patterns are essentially the same from what we documented previously. Demand for higher educated workers tended to grow or remained flat during the 1990s, but the trend started to reverse by the new millennium.

Figure 6.2: Demand Trends Using Alternative Model Specification

(a) Demand Trends Alternative Model: High School/Primary ($\log \beta_t$)

(b) Demand Trends Alternative Model: Skilled/Unskilled ($\log \alpha_t$)

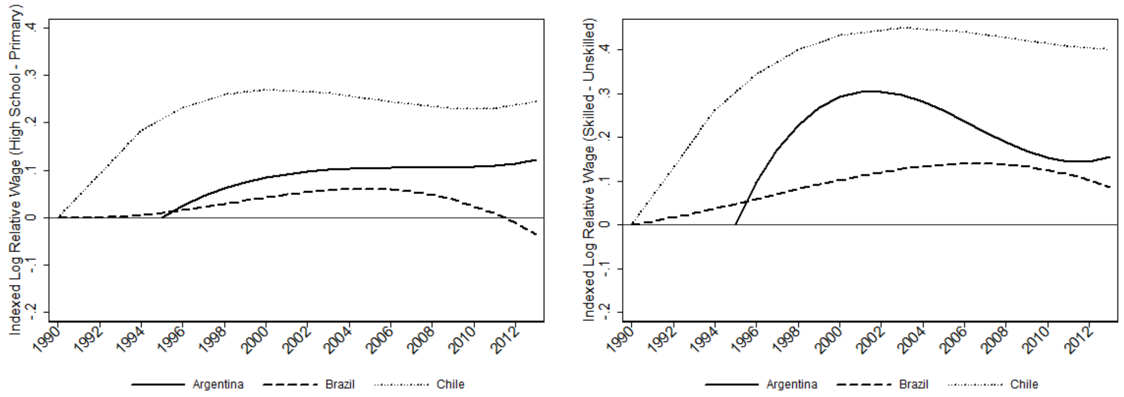

Notes: the Panel (a) depicts the estimated relative demand trends between workers with at most a high school degree and those with only primary education. the panel (b) depicts the estimated relative demand trends between skilled and unskilled workers. These relative demands correspond to the first and last stage of the alternative model discussed in Appendix A.3. Each series is scaled so that it takes a value of zero at the first year in which data for the country is available.

7 Conclusions

After a decade of stagnant or rising earnings inequality, the distance between top and bottom earners in Latin America fell sharply during the 2000s. This trend was in sharp contrast to the experience of developed countries around the same period. This paper has offered a detailed accounting of the main factors behind the evolution of earnings inequality in three of the largest countries in the region: Argentina, Brazil and Chile.

The first suspect for changes in the wage structure are changes in the composition of the labor force. The three countries studied here were subject to similar employment changes, although in varying degrees: employees are aging, becoming more educated and are more likely to be females. We construct counter-factual wage distributions where the returns to labor market characteristics are kept fixed to evaluate how these changes in the composition of employment may have affected the distribution of wages. Our results are unambiguous. Changes in the composition, in particular, the educational upgrading, were inequality enhancing. Hence, while composition changes may have contributed to increasing wage inequality before the 2000s, they cannot explain the substantial decline of the last decade.

The decomposition also allows us to build counterfactual distributions where composition changes are kept constant to evaluate the role played by changes in the returns to labor market attributes. The analysis suggests a distinct role of the education and experience premiums. The decline of the experience premium is key to explain reductions of upper-tail (90/50 earnings gap) inequality. This is because reductions of the experience premium were stronger among the highly educated, perhaps reflecting some skill obsolescence. In contrast, a falling schooling premium bears a much higher weight in the reduction of inequality below the median (50/10 earnings ratio). This was driven by a much faster reduction of the high school premium vis-à-vis workers who have at most completed primary education than the reduction of the college premium.

To link changes in the schooling and experience premiums to the observed changes in labor supply we have built a nested CES model where there is imperfect substitution across experience and education groups. A combination of a relative trend in demand that favoured high school educated vis-à-vis primary educated workers during the 1900s, but slightly reversed during the 2000s, and a rapid educational upgrade go a long way in explaining the high school/primary schooling premiums. The story in the case of college educated workers is similar. Rising supply of college educated workers has pushed the college premium downwards, but this is not enough to explain the reduction that started around 2000. We find instead that the demand for college educated workers fell during the 2000s in the three countries. Changes in the experience premium have also responded to the evolution of relative supplies, suggesting imperfect substitutability (our estimate of the elasticity of substitution is between 3 and 9). However, as happened with the college educated workers, the changing experience profile of the labor force is not sufficient to explain the reduction of the returns that took place during the 2000s.

We show that expanding our empirical model to account for changes in the minimum wage and cyclical conditions of the labor market does not have major impact on our results. The sharp increases of the minimum wages in these countries have contributed to the decline of the college premiums, but the importance of the relative supply channel, and the estimated patterns of relative demands, are mostly unaffected once the model is extended to account for this fact. The results are also robust to using different measures of labor supply, and to alternative specifications of the underlying theoretical model.

We have focused our exposition on the common factors that have driven earnings inequality in the three countries, but in spite of commonalities, substantial heterogeneity remains. Brazil witnessed a much more pronounced reduction in earnings inequality than Argentina and Chile. Top and bottom inequality reductions took place in parallel in Brazil and Chile, while the decline of inequality in Argentina was fundamentally driven by the evolution in the upper half of the distribution. In Chile the demand for high skilled workers increased much more rapidly during the 1990s than in Brazil and Argentina, and in spite of the recent fall remains above its 1990 level by 2013. Understanding the sources of this heterogeneity constitutes an important avenue for further research.

A Appendix

A.1 Data and Variable Construction

The household surveys used in Argentina for the period between 1995 and 2003 is the Encuesta Permanente de Hogares (EPH), collected by the Instituto Nacional de Estadística (INDEC). This survey was replaced by the Encuesta Permanente de Hogares Continiua (EPH-C) after 2003, breaking the series. The transition between the EPH and the EPH-C included changes in the questionnaires and the frequency in which the surveys were collected, so results should be interpreted with this caveat in mind. The geographical coverage was also extended to include additional agglomerates. In order to maintain a consistency over the period of study we only use the agglomerates that are present in both surveys. The EPH and the EPH-C are representative for urban areas, but close to 90% of the population in Argentina live in urban centers.

The survey used in Brazil is the Pesquisa Nacional por Amostra de Domicilios (PNAD), collected by the Instituto Brasilero de Geografía y Estadísticas (IBGE). The PNAD is a nationally representative survey that has been carried out on a yearly basis since 1967. We use the different waves starting from the year 1990. Due to exceptional circumstances the survey was not collected in 1994 and 2000.

The household survey used for Chile is the Encuesta de Caracterización Socioeconómica Nacional (CASEN). The CASEN is a nationally representative household survey carried out by the Ministry of Planning through the Department of Economics at the Universidad de Chile. The survey was first implemented in 1987, and was carried out every two years from 1990 to 2000, and every three years thereafter. We use the different waves from 1990 to 2013.

We constructed variables capturing the educational attainment and potential experience of all individuals in the sample. Although the countries we analyse differ in the structure of their educational systems, the SEDLAC project has attempted to homogenize the information from the different countries to make it comparable.[21] We use SEDLAC's coding in the construction of the educational attainment series. In particular, we define 5 possible levels of educational attainment: (a) primary education completed or less; (b) high school incomplete; (c) high school completed; (d) college incomplete; and (e) college completed or more. Potential experience is defined as the result of subtracting the total number of years of education completed (plus 6) to the age of the individual.

Although we define 5 possible levels of education attainment, we mostly work with three categories: primary or less, high school completed and college completed. Individuals

[21] See CEDLAS and The World Bank (2014) for a detailed description of the SEDLAC database

with incomplete levels of education are distributed equally between the previous and next completed level. For example, mean real hourly earnings of workers with college education are calculated as a weighted average between the observed mean wages of this group and the observed mean wages of workers with college incomplete. The weight of the latter group is equal to half of their actual number. This also implies that in the labor supplies used in the model, the supply of workers with primary education completed or less include half of the total supply of workers of the high school incomplete category. The supply of workers with high school education completed include both half of the supply of workers with high school incomplete and half of the supply of workers with college incomplete. Finally, the supply of workers with college education completed includes half of the total supply of workers with college incomplete.

Each survey includes a question asking workers for the total monetary income from labor in a reference period. This is the variable that we use throughout the paper to capture labor earnings. The variable is divided by the self reported total number of hours worked to obtain hourly earnings. The series are converted into real terms using the consumer price index of the respective countries.[22] We have restricted the sample to individuals between the ages of 16 and 65, and only use earnings of full time workers (individuals that self reported working for more than 35 hours in the reference week).

The composition adjusted earnings of aggregate groups are constructed using a fix-weighted average of the different sex-education-experience sub-groups. We first run a regression of log hourly earnings on the full set of covariates that include indicators for the 5 education categories, 7 dummies for potential experience in five year intervals, and all possible interactions. The regression is estimated separately for males and females in each available country-year. The predicted log wages from these regressions are evaluated for the 70 sub-groups, and a weighted average is estimated when aggregating to broader groups. The weights are equal to the mean employment share of each sub-group across all years.

A.2 Using RIF to Decompose Changes in Distributional Statistics Beyond the Mean

Firpo et al. (2007, 2009) propose a methodology that allows extending the traditional Oaxaca-Blinder decomposition to distributional statistics beyond the mean. This is achieved through the use of influence functions (IF). Influence functions measure the effect that an infinitesimal amount of "errors" have on a given estimator (Cowell and Victoria-Feser, 1996), but they also have properties that allows us to model the sensitivity of a given

[22] Due to inconsistencies found in the official Consumer Price Index in Argentina (see Cavallo (2013)), we use the information from PriceStats (`http://www.statestreet.com/ideas/pricestats.html`) to deflate nominal wages in this country.

unconditional wage quantile to a change in a set of covariates. To see this, let $q_\tau(F_W)$ be τth quantile of the distribution of wages, expressed in terms of the cumulative distribution $F_W(w)$. Define the following mixture distribution:

$$G_{W,\epsilon} = (1-\epsilon)F_W + \epsilon H_W \quad for \quad 0 \leq \epsilon \leq 1 \tag{A.1}$$

where H_W is some perturbation distribution that only puts mass at the value w. In that case, $G_{W,\epsilon}$ is a distribution where, with probability $(1-\epsilon)$, the observation is generated by F_W, and with probability ϵ, the observation takes the arbitrary value of the perturbation distribution. By definition, the influence function corresponds to:

$$IF(w; q_\tau, F_W) = lim_{\epsilon \to 0} \ \frac{q_\tau(G_{W,\epsilon}) - q_\tau(F_W)}{\epsilon} \tag{A.2}$$

where the expression is analogous to the directional derivative of q_τ in the direction of H_W. Analytical expressions for influence functions have been derived for many distributional statistics.[23] The influence function in the case of the τth quantile takes the form:

$$IF(w; q_\tau, F_W) = \frac{\tau - \mathbb{1}[w \leq q_\tau]}{f_W(q_\tau)} \tag{A.3}$$

Where $\mathbb{1}[\cdot]$ is an indicator function and f_W is the PDF.[24] Using some of the properties of influence functions, a direct link with the traditional Oaxaca-Blinder approach can be established. In particular, a property that is shared by influence functions is that, by definition, the expectation is equal to zero.

$$\int_{-\infty}^{+\infty} IF(w; q_\tau, F_W) dF(w) = 0 \tag{A.4}$$

Firpo et al. (2009) propose a simple modification in which the quantile is added back to the influence function, resulting in what the authors call the Recentered Influence Function (RIF).

$$RIF(w; q_\tau, F_W) = q_\tau + IF(w; q_\tau, F_W) \tag{A.5}$$

[23] Essama-Nssah and Lambert (2011) provides a comprehensive list of influence functions for different distributional statistics.

[24] Note that the influence function in this case depends on the density. In order to obtain the empirical density the authors propose non-parametric kernel density estimation.

The importance of this transformation lies in the fact that the expectation of the RIF is precisely the quantile q_τ. With this result, Firpo et al. (2009) show that we can model the conditional expectation of the RIF as a linear function of the explanatory variables.

$$E[RIF(w_t; q_\tau, F_{W,t}|X_t)] = X_t'\gamma_t \tag{A.6}$$

Moreover, if we apply the law of iterated expectations to equation A.6, the end result is an expression that directly relates the impact of changes in the expected values of the covariates on the unconditional quantile q_τ. Note that this result is all that is required to extend the Oaxaca-Blinder decomposition to quantiles, since the basic components of the method are all present in equation (A.6).

Estimation of equation (A.6) can be done by OLS, and only requires replacing the dependent variable, $\log w_t$, in the original wage setting model, with the RIF of the quantile q_τ. The interpretation of the estimates $\widehat{\gamma}_t$ can be thought of as the effect of a small change in the distribution of X on q_τ, or as linear approximation of the effect of large changes of X on q_τ (Firpo et al., 2007).

A.3 Alternative Model Specification

In this section we present an alternative specification of the production function of the model in Section 4.1. We mostly follow the work of Manacorda et al. (2010), with some small modifications to allow for comparability with our baseline results. Production in the economy is also modelled using a nested constant elasticity of substitution (CES) function with three levels. The first level is identical to the one we use

$$Y_t = \lambda_t \left(L_{Ut}^{\rho} + \alpha_t L_{St}^{\rho} \right)^{1/\rho} \tag{A.7}$$

with the parameters having the same interpretation. In the second level, the labor from skilled and unskilled workers is divided into 7 potential experience sub-groups, aggregating them with a productivity-weighted CES combination of the form

$$L_{Mt} = \left(\sum_{A=1}^{7} \phi_{MA} L_{MAt}^{\theta} \right)^{1/\theta} \quad \text{for} \quad M = S, U \tag{A.8}$$

where A indexes the potential experience groups; ϕ_{MA} is a time-invariant parameter capturing differences in relative productivities between potential experience groups; and θ is a function of the elasticity of substitution: $\sigma_\theta = \frac{1}{1-\theta}$. Two key differences are worth pointing out. First, the second level of the production function aggregates labor by experience, not by skill sub-group. The ordering between the second and third levels is then shifted. Second, the model assumes that there are no relative demand/productivity changes between workers with different levels of potential experience. This follows from the assumption that the respective parameters (ϕ_{MA}) are time-invariant, which largely simplifies the estimation.

Finally, the supply of labor from workers with a given potential experience within the unskilled group is composed of a CES combination of labor from the two lower schooling levels

$$L_{UAt} = \left(L^{\delta}_{PAt} + \beta_t L^{\delta}_{HAt}\right)^{1/\delta} \tag{A.9}$$

where P and H denote workers with primary education or less and high school completed respectively; β_t is a time-variant measure of the relative productivity between the two low education levels; and δ is a function of the elasticity of substitution between the two groups: $\sigma_\delta = \frac{1}{1-\delta}$. Note that β_t is constant across potential experience groups, so relative demand shifts are common in this dimension. Finally, the natural logarithm of the two time-variant parameters (α_t and β_t) are allowed to move according to cubic time trend.

Two differentiating factors between this specification and the work of Manacorda et al. (2010) are worth pointing out. First, we allow for differential demand trends within low skilled workers, which they assume to be constant. Second, we allow for a more flexible specification of the demand trends by fitting a cubic polynomial instead of a linear time trend. The fact that we observed some sharp trend reversals in our baseline model is indicative that this difference could be quantitatively important.

References

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In Ashenfelter, O. and Card, D. (eds), *Handbook of Labor Economics*. Elsevier, *vol. 4B*, chap. 12.

Alvaredo, F. and Gasparini, L. (2015). Chapter 9 - recent trends in inequality and poverty in developing countries. In Atkinson, A. B. and Bourguignon, F. (eds), *Handbook of Income Distribution*. Elsevier, Handbook of Income Distribution *2*, 697 – 805.

Autor, D., Katz, L. and Kearney, M. (2005). Rising wage inequality: The role of composition and prices. NBER Working Paper Series 11628.

Autor, D., Katz, L. and Kearney, M. (2008). Trends in U.S. inequality: Revising the revisionists. *The Review of Economics and Statistics* 90(2): pp. 300–323.

Azevedo, J. P., Inchauste, G. and Sanfelice, V. (2013). Decomposing the recent inequality decline in Latin America. Policy Research Working Paper Series 6715, The World Bank.

Behrman, J., Birdsall, N. and Szekely, M. (2007). Economic Policy Changes and Wage Inequality in Latin America. *Economic Development and Cultural Change* 56(1): pp. 57–97.

Bourguignon, F., Ferreira, F. H. and Lustig, N. (2005). *The Microeconomics of Income Distribution Dynamics in East Asia and Latin America*. No. 14844 in World Bank Publications. The World Bank.

Card, D. and Lemieux, T. (2001). Can falling supply explain the rising return to college for younger men? a cohort-based analysis. *The Quarterly Journal of Economics* 116(2): pp. 705–746.

Cavallo, A. (2013). Online and Official Price Indexes: Measuring Argentina's Inflation. *Journal of Monetary Economics* 60(2): pp. 152–165.

CEDLAS and The World Bank (2014). A Guide to SEDLAC: Socio-Economic Database for Latin America and the Caribbean.

Cornia, A. (2010). Income Distribution Under Latin America's New Left Regimes. *Journal of Human Development and Capabilities* 11(1): pp. 85–114.

Cowell, F. and Victoria-Feser, M.-P. (1996). Robustness properties of inequality measures. *Econometrica* 64(1): pp. 77–101.

Cragg, M. I. and Epelbaum, M. (1996). Why has wage dispersion grown in Mexico? is it the incidence of reforms or the growing demand for skills? *Journal of Development Economics* 51(1): pp. 99–116.

Erten, B. and Ocampo, J. A. (2013). Super Cycles of Commodity Prices Since the Mid-Nineteenth Century. *World Development* 44: 14–30.

Essama-Nssah, B. and Lambert, P. (2011). Influence functions for distributional statistics. ECINE Working Paper 236.

Ferreira, F., Leite, P. and Litchfield, J. (2008). The Rise and Fall of Brazilian Inequality. *Macroeconomic Dynamics* 12(2): pp. 199–230.

Firpo, S., Fortin, N. and Lemieux, T. (2007). Decomposing wage distributions using re-centered influence functions. Unpublished manuscript, PUC-Rio and UBC.

Firpo, S., Fortin, N. and Lemieux, T. (2009). Unconditional Quantile Regressions. *Econometrica* 77(3): pp. 953–973.

Gasparini, L., Cruces, G. and Tornarolli, L. (2011). Recent Trends in Income Inequality in Latin America. *Economía* 11(2): pp. 147–201.

Gasparini, L. and Lustig, N. (2011). The Rise and Fall of Income Inequality in Latin America. In Ocampo, J. A. and Ros, J. (eds), *The Oxford Handbook of Latin American Economics*. Oxford University Press.

Johnson, M. and Keane, M. (2013). A Dynamic Equilibrium Model of the US Wage Structure, 1969-1996. *Journal of Labour Economics* 31(1): pp. 1–49.

Kahhat, J. (2010). Labor Earnings Inequality: The Demand for and Supply of Skills. In López-Calva, L. F. and Lustig, N. (eds), *Declining Inequality in Latin America: A Decade of Progress?*. Brookings Institution Press, chap. 2.

Katz, L. and Autor, D. (1999). Changes in the Wage Structure and Earnings Inequality. In Ashenfelter, O. and Card, D. (eds), *Handbook of Labor Economics*. Amsterdam: North-Holland, *vol. 3A*, chap. 26.

Katz, L. and Murphy, K. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics* 107(1): pp. 35–78.

Lemieux, T., MacLeod, B. and Parent, D. (2009). Performance Pay and Wage Inequality. *Quarterly Journal of Economics* 124(1): pp. 1–49.

Levy, S. and Schady, N. (2013). Latin america's social policy challenge: Education, social insurance, redistribution. *Journal of Economic Perspectives* 27: 193–218.

Londoño, J. L. and Szekely, M. (2000). Persistent Poverty and Excess Inequality: Latin America, 1970-1995. *Journal of Applied Economics* 3(1): pp. 93–134.

Lopez-Calva, L. F. and Lustig, N. (2010). *Declining Inequality in Latin America: A Decade of Progress?*. Brookings Institution Press.

López-Calva, L. F. and Lustig, N. (2010). Explaining the decline in inequality in Latin America: Technological change, educational upgrading, and democracy. In López-Calva, L. F. and Lustig, N. (eds), *Declining Inequality in Latin America: A Decade of Progress?*. Brookings Institution Press, chap. 1.

Lustig, N., Lopez-Calva, L. and Ortiz-Juarez, E. (2013). Deconstructing the Decline in Inequality in Latin America. Tulane Economic Working Paper Series.

Manacorda, M., Sánchez-Paramo, C. and Schady, N. (2010). Changes in Returns to Education in Latin America: The Role of Demand and Supply of Skills. *Industrial and Labor Relations Review* 63(2): pp. 307–326.

Murphy, K. and Welch, F. (1992). The Structure of Wages. *The Quarterly Journal of Economis* 107(1): pp. 285–326.

Sánchez-Páramo, C. and Schady, N. (2003). Off and Running? Technology, Trade, and the Rising Demand for Skilled Workers in Latin America. The World Bank. Policy Research Working Paper 3015.

The World Bank (2016). The Commodity Cycle in Latin America: Mirages and Dilemas. Semiannual Report, Office of the Regional Chief Economist.