

Skill Acquisition and the Dynamics of Trade-Induced Inequality*

Eliav Danziger[†]
Princeton University

Job Market Paper

Latest version: http://scholar.princeton.edu/files/Danziger_JMP

January 6, 2014

Abstract

This paper quantifies the impact of trade liberalization on wage inequality between workers of different skill levels and across age groups. I propose a model in which trade liberalization increases the demand for skill due to production share reallocation across firms and technology switching. However, unlike in the existing literature, I endogenize the skill supply by supplementing the skill-demand side of the model with an overlapping-generations model of skill acquisition. I calibrate the model to 2007 US data and simulate the economy's transition path in response to the removal of policy trade barriers. Workers have rational expectations and, therefore, must take into account the general-equilibrium effects on wages of changes in skill supply during the economy's transition. I find that the aggregate gains from trade liberalization, defined as the increase in discounted real earnings relative to their pre-liberalization level, are 5.9%. However, these gains are not distributed evenly among workers. For those alive at the time of implementation of the new trade policy, the oldest educated workers' discounted real lifetime earnings increase by 9.9%, while the oldest uneducated workers' discounted real lifetime earnings increase by only 1.5%. On the one hand, ignoring the economy's transition leads to an understatement of trade-induced inequality as this fails to account for transitory inequality. On the other hand, ignoring the endogeneity of the skill supply leads to an overstatement of trade-induced inequality as this fails to account for the equalizing effect of the endogenous skill-supply adjustment.

Keywords: Trade liberalization; Wage inequality; Skill premium; Education

JEL Classifications: C68; F16; F66; I24; J24; J31

*I am grateful to Esteban Rossi-Hansberg, Stephen Redding and Gene Grossman for their invaluable guidance and support. I thank Oleg Itskhoki, Greg Kaplan, Eduardo Morales, Richard Rogerson, Felix Tintelnot and Jonathan Vogel for their insightful comments and suggestions. Financial support from the International Economics Section at Princeton University is greatly appreciated.

[†]Department of Economics, Princeton University, Princeton, NJ 08544, USA. Homepage: <http://scholar.princeton.edu/edanzige>. Email: edanzige@princeton.edu.

1 Introduction

The main objective of this paper is to quantify the differential impact of trade liberalization on workers of different education levels and across age groups when the supply of educated workers is endogenous. The model I develop to facilitate the quantitative analysis is motivated by two insights that have generally been overlooked in the literature. First, if trade induces a change in the relative demand for educated workers, then workers will respond by adjusting their schooling decisions. Second, the adjustment of the supply of educated workers to a change in the relative demand for their services is a protracted process.

These insights are not merely theoretical. Mexico provides a case in point. Beginning in the mid-1980's Mexico implemented a series of trade-liberalizing policies that culminated in the implementation of NAFTA in 1994. During this period of liberalization, increased demand for college graduates led to a steady increase in the college premium. However, after 1994 the college premium experienced a prolonged decline with a contemporaneous increase in the supply of college graduates.¹ A similar pattern has been documented in Korea where, after trade liberalizations and the implementation of export-promoting policies, the income gap between college graduates and high-school graduates first widened from 1971 to 1976 and subsequently narrowed in response to a growing supply of college graduates.²

These episodes are the empirical manifestation of the two theoretical insights that motivate this paper. In both cases, an increase in the demand for college graduates and the concomitant increase in the college premium were followed by a long decline in the college premium accompanied by an augmentation of the supply of college graduates. These episodes suggest that skill-supply adjustment is an important margin in accounting for the effects of trade-induced shifts in relative skill demand. Furthermore, these episodes demonstrate that skill adjustment can be a drawn-out process. Thus, ignoring the endogeneity of the skill supply and considering only long-run impacts provides, at best, only a partial picture of the effects of trade-induced shifts in relative skill demand.

In order to incorporate these observations into the quantitative analysis of the impact of trade liberalization on inequality, I build an overlapping-generations model with endogenous skill acqui-

¹See Robertson (2007) and Campos-Vazquez (2010) who attribute the decrease in the college premium to an increase in the supply of college graduates after NAFTA.

²See Kwark and Rhee (1993) and Kim and Topel (1995).

sition. To get a fuller picture of the effects of trade liberalization, I consider the economy's entire transition path following liberalization. This contrasts with the existing literature on trade and inequality in which skill supplies are generally taken to be exogenous and only steady-state equilibria are considered.

Quantifying the dynamic effects of trade with endogenous education decisions introduces significant technical challenges into the analysis. The reason is that, because workers have rational expectations, present education decisions of workers depend on future wages. However, future wages respond to present and future education decisions through the general-equilibrium impact of skill supplies on the skill premium. The evolution of the economy, therefore, consists of a path of interdependent endogenous variables rather than a series of independent static equilibria. As a consequence, the entire equilibrium path must be solved by employing iterative numerical techniques that are computationally intensive.³

Despite these challenges, the model I develop is sufficiently tractable to allow for a solution of the entire equilibrium transition path. At the same time, the model remains rich enough to provide meaningful predictions about the impact of trade liberalization on workers along several dimensions. In particular, the model provides predictions about the impact of trade liberalization on workers that differ in educational attainment and age, as well as in birth cohort.

In the model, uneducated workers can either work full time and supply low-skill labor or pursue an education that ultimately will enable them to provide high-skill labor. Pursuit of an education, however, is a costly activity in terms of both tuition and time. Complementarity between education and ability implies that only the more able individuals find it worthwhile to invest in an education. Moreover, since workers have a finite expected lifetime, older workers are less likely to pursue an education as they stand to reap the fruits of their investment for a shorter period of time.

An increase in relative skill demand, which in the model is triggered by trade liberalization, drives up the returns to education at impact. The high returns to education mean that education now becomes a good investment for some workers for whom previously it was not. This precipitates an increase in workers seeking an education that, over time, augments the skill supply and thereby

³Other analyses of the dynamic impacts of trade on workers have sought to circumvent these difficulties in various ways. For example, Falvey et al. (2010) eliminate the general-equilibrium effects of present decisions on future wages by studying a small open economy. In such a case, wages depend only on exogenously given world prices.

depresses the high returns to education.

Skill supplies do not adjust immediately to the increase in relative skill demand for two reasons. First, pursuing an education is a time-consuming activity. Second, old workers, who would have acquired an education had post-liberalization conditions prevailed in their youth, no longer find it profitable to do so. Skill supplies can only fully adjust when these workers are replaced in the labor force by younger cohorts. The rate at which skill supplies adjust determines, through its general-equilibrium effect on the skill premium, the differential impact of trade liberalization on workers during the economy's transition phase.

Trade liberalization does not impact all workers symmetrically. Educated and uneducated workers are affected differentially by liberalization because of the asymmetric effects on the relative demand for their services. The dynamic effects on these workers depend on the trajectory of the skill premium during the transition. The implications of trade liberalization for old and young differ for two reasons. First, old uneducated workers are less likely than young workers to acquire an education in response to the increased returns to education. Second, old workers' lifetime earnings are affected only by the wages in the near future, whereas young workers' lifetime earnings depend also on wages in the more distant future.

The catalyst for the increase in the returns to education and the associated skill-supply adjustment is the change in relative skill demand induced by trade liberalization. The trade literature has proposed several mechanisms through which trade can induce shifts in relative skill demand. Traditionally, the literature has emphasized intersectoral shifts in labor demand as the key link between trade and skill demand.⁴ Recently, a growing literature has emphasized the role of firm heterogeneity and market-share reallocations toward more efficient firms within a given sector as the driving force behind trade-induced shifts in relative skill demand.⁵

In the model, I focus on shifts in relative skill demand emanating from between-firm labor reallocations. This is in line with the empirical finding that trade shocks lead to substantial labor reallocation between firms within sectors, but to little reallocation across sectors.⁶ In addition, Burstein and Vogel (2011) find that, of the increase in the skill premium associated with increased

⁴See Goldberg and Pavcnik (2007) for a survey on the link between trade and inequality in developing countries. The main focus in that survey is on older trade theories in the spirit of Heckscher-Ohlin.

⁵See Tybout (2000) for an early survey of the evidence on the relationship between firm heterogeneity and trade.

⁶See, for example, Haltiwanger et al. (2004) and Wacziarg and Wallack (2004).

trade in the United States, two-thirds are attributable to within-sector reallocations, while only one third is attributable to between-sector reallocations.

To rationalize trade-induced shifts in relative skill demand, the firm side of the model builds on Bustos (2011) in which heterogeneous firms choose a production technology from a menu of technologies that differ in productivity and skill intensity. Trade liberalization induces the reallocation of production shares towards exporters as in Melitz (2003), which affects relative skill demand through two channels. First, if exporters, on average, employ more skill-intensive technologies than non-exporters, this reallocation will tend to increase relative skill demand. Second, exporters expand their production and therefore upgrade their technologies, while the reverse is true for non-exporters. To the extent that technologies differ in their skill intensity, this technology switching shifts relative skill demand.

I calibrate the model to 2007 US data and simulate the economy's response to a once-and-for-all removal of policy trade barriers. The gains from trade, defined as the increase in discounted real lifetime earnings relative to their pre-liberalization level, aggregated over present and future generations, are 5.9%. However, the gains generated by the policy are not evenly divided among workers. Old educated workers alive at the time of implementation of the new trade policy gain 9.9%, making them the biggest winners from trade. In contrast, old uneducated workers gain only 1.5%. In general, older educated workers gain more than younger educated workers, while the opposite is true for uneducated workers. However, every educated worker gains more than any uneducated worker.

I find that there is a large increase in inequality immediately following trade liberalization, both between educated and uneducated workers and across age groups. However, over time, the skill-supply adjustment mitigates the adverse distributional impact of trade liberalization. Nevertheless, even in the long run, trade liberalization does lead to a small increase in inequality.

Finally, I assess the importance of accounting for the endogeneity of the skill supply and the economy's transition path. On the one hand, I find that assuming that the skill supply is exogenous leads to a substantial overstatement of trade-induced inequality as it does not account for the equalizing effect of the endogenous skill-supply adjustment. On the other hand, ignoring the economy's transition and considering instead only steady-state equilibria understates trade-induced inequality.

Indeed, much of the inequality is manifested during the transition during which skill supplies are adjusting to the increased relative skill demand.

The remainder of the paper is organized as follows: I begin by briefly outlining the relationship of this paper to the existing literature. Section 2 lays out the theoretical model. Section 3 describes the calibration of the model's parameters to 2007 US data. Section 4 describes the simulation of the calibrated model's response to a counterfactual trade liberalization and reports the results. Section 5 concludes.

Relation to Literature

This paper contributes to a growing literature emphasizing the role of firm heterogeneity as a crucial link between trade and inequality.⁷ Bernard and Jensen (1997) provide the first empirical evidence of this link. They show that a substantial fraction of the increase in the skill premium observed in the 1980's occurred through labor reallocations between plants within industries, and, in particular, the effect occurred almost entirely among exporting firms. The theory seized on these empirical findings beginning with Melitz (2003). The key theoretical insight in that model is that trade shifts production shares from less-efficient to more-efficient firms. However, with homogenous workers and competitive wage setting, this mechanism by itself cannot explain the link between firm heterogeneity and trade-induced inequality.

The literature has proposed several mechanisms to explain why more-efficient firms pay similar workers different wages than less-efficient firms, thereby linking trade and inequality. Egger and Kreckemeier (2009) posit that more-efficient firms pay higher wages because of fair-wage considerations, which leads to increased inequality. Davis and Harrigan (2007), in an efficiency-wage model, show that inequality can either increase or decrease in response to trade depending on the correlation between a firm's efficiency in production and its efficiency in monitoring. Helpman, Itskhoki and Redding (2010) add search frictions, bargaining and screening to a Melitz-style model to explain why more-productive firms pay higher wages and hence how production-share reallocation toward those firms can increase inequality. In these models, workers are ex-ante homogeneous and therefore

⁷A consumer side approach to linking trade with inequality has recently been suggested by Fajgelbaum, Grossman and Helpman (2011).

non-competitive wage setting is needed to rationalize inequality.⁸

In contrast to the aforementioned models, in my model, wage setting is competitive, but workers command different wages depending on whether or not they are educated. This is related to the approach in Burstein and Vogel (2011) who study a model with competitive wage setting and skilled and unskilled labor. In their model, high skill intensity in production is an innate characteristic of more-efficient firms. As such, the reallocation of production shares toward these firms drives up the relative demand for skilled workers. In contrast, in my approach, firms choose from a menu of technologies which differ in skill intensity. Thus, skill intensity is a property of the technology and not of the firm. Therefore, in addition to the production-share reallocation effect, trade influences relative skill demand because it induces firms to switch technologies.⁹

My approach to linking relative skill demand to trade is most closely related to Bustos (2011). The key differences are that I expand the technology menu to more than two choices, and, more importantly, I do not assume that more-expensive and more-productive technologies are also more skill intensive. Indeed, with this assumption, the increase in relative skill demand caused by a reduction in trade costs would be hardwired into the model. In the calibration of the theoretical model, I find that, in fact, more-productive technologies are not necessarily more skill intensive, but rather that there exists an inverted U-shape relationship between a technology's productivity and its skill intensity.

This paper is also related to a nascent literature on the dynamic impact of trade on workers.¹⁰ One strand of this literature has highlighted frictions in intersectoral mobility of workers as the source of slow adjustment to trade liberalization. Artuç, Chaudhuri and McLaren (2010) develop a dynamic model in which workers do not immediately move to a sector even if it offers higher wages than their current sector because of mobility costs. Artuç (2009) and Dix-Carneiro (2013)

⁸Verhoogen (2008) suggests that in developing countries exporters produce high-quality goods. To the extent that the production of high-quality goods is more skill-intensive than the production of low-quality goods, an expansion of export markets will increase relative demand for skills. Matsuyama (2007) introduces a model in which serving the export market is a more skill-intensive activity than serving the domestic market as it requires knowledge of international business as well as language skills. See Harrison, McLaren and McMillan (2011) for a review of the recent theoretical and empirical contributions to the literature on trade and inequality.

⁹Lileeva and Trefler (2010) show that Canadian firms that started exporting after trade liberalization upgraded technology. Bustos (2007) provides evidence that Argentinean firms induced to export by the implementation of MERCOSUR upgraded technology. Bustos (2011) shows that this technology upgrade was associated with an increase in skill intensity among upgraders, suggesting that these firms upgraded to more skill-intensive technologies.

¹⁰There has also been some recent research into the dynamic impact of trade liberalization on firms in the presence of sunk costs. See, for example, Atkeson and Burstein (2007) and Costantini and Melitz (2007).

extend this model to incorporate sector-specific human capital as an added barrier to intersectoral mobility.¹¹ In my approach, adjustment to trade liberalization is slow because it takes time for skill supplies to adjust to changes in relative skill demand.

Given the emphasis in the trade literature on the effect of trade on the skill premium, there has been surprisingly little research on the effect of trade on skill acquisition. Early work by Findlay and Kierzkowski (1983) analyzes the consequences of introducing endogenous skill acquisition into a Heckscher-Ohlin model. More recent research has modified this approach and extended it to a dynamic setting. Harris and Robertson (2013) embed a dynamic model of skill acquisition into a small open-economy trade model. They use their model to study the potential effect of globalization on India and China. Their model differs from mine in focus inasmuch as they study the effect of trade on developing countries, whereas my focus is on the United States. In addition, in their model households are homogeneous and decide how much of their time to allocate to skill acquisition. Thus, although their model offers predictions about the effect of skill acquisition on the skill premium, the relationship between the skill premium and inequality remains unclear.

Falvey, Greenaway and Silva (2010) study a dynamic model of skill acquisition in a small open economy in which workers differ in ability and age. Their model is related to mine in that it sheds light on the differential impact of trade on workers by ability and age. A key difference between their work and mine is that the small open-economy setting ensures that wages are exogenously given by world prices, and, as such, there is no general-equilibrium feedback between workers' decisions and future wages. Indeed, a defining feature of my framework is that tomorrow's wages are determined by today's schooling decisions and vice versa.

2 The Model

2.1 Setup

The model is set in discrete time with periods indexed by t . There are two identical countries, Home and Foreign, indexed by $i \in \{H, F\}$. Each economy is populated by workers who maximize expected discounted consumption of a nontradeable final good that is a constant elasticity of substitution

¹¹In a related paper Cosar (2010) merges sector-specific human capital with search frictions to determine the relative importance of those two factors in the labor market's adjustment to trade liberalization.

(CES) aggregate of intermediate good varieties with elasticity σ . Varieties of the intermediate good are produced by profit-maximizing firms that hire the labor services of workers and also make use of the final good to cover fixed production costs. In any period t , the supply of the final good in country i is

$$Q_{it} = \left[\int_{\omega \in \Omega_{it}} q_{it}(\omega)^\rho d\omega \right]^{1/\rho},$$

where Ω_{it} is the endogenous set of intermediate varieties available in country i in period t , $q_{it}(\omega)$ is the quantity supplied of a given intermediate variety ω and $0 < \rho = (\sigma - 1) / \sigma < 1$. Given this CES structure, the price of the final good is

$$P_{it} = \left[\int_{\omega \in \Omega_{it}} p_{it}(\omega)^{1-\sigma} d\omega \right]^{1/(1-\sigma)},$$

where $p_{it}(\omega)$ is the consumer price of intermediate variety ω in country i in period t . The price of the final good will henceforth be normalized to unity which means that in each period t and each country i final good expenditure R_{it} is equal to the supply of the final good Q_{it} .

For simplicity, the exposition will focus on symmetric equilibria.¹² The model and subsequent results are therefore presented from the point of view of the Home country, with the understanding that the results are identical for the Foreign country.

2.2 Firms

In this section, I describe the firm side of the model, which is an extension and modification of Bustos (2011). There is a mass, M , of infinitely-lived, profit-maximizing intermediate-good producing firms in the Home country. Each firm has the ability and sole right to produce a unique variety of the intermediate good. Firms are owned equally by all the workers in the country, so that profits are shared equally among all workers.¹³

In each period a firm must choose a technology, $v \in \{0, 1, \dots, V\}$, from a menu of available technologies. Each technology is characterized by a triplet $(\alpha_v, \lambda_v, f_v)$, where α_v is a Hicks-neutral

¹²Although I do not prove analytically that there exist only symmetric equilibria, every simulation converges to a symmetric equilibrium. Every simulation converges to a symmetric equilibrium even though I do not impose symmetry and even when I choose different initial guesses for the wage path for each country. For more details on the simulations see Appendix C.

¹³This is the Chaney (2007) variant of a Melitz model.

productivity factor, λ_v is the skill intensity of the technology and f_v is the fixed cost of using technology v for one period in terms of the final good. Importantly, unlike in Bustos (2011), I allow the technology menu to include an arbitrary number of technologies rather than just two, and I do not take a stand on the relationship between $(\alpha_v, \lambda_v, f_v)$ for different technologies.

The production function of a firm employing technology v is

$$q(\ell, h) = \varphi \alpha_v \min \left\{ \ell, \frac{h}{\lambda_v} \right\},$$

where ℓ and h are the measures of low- and high-skill labor, respectively, hired by the firm, and φ is a firm-specific idiosyncratic productivity factor independent of the chosen technology. The idiosyncratic productivities, φ , are distributed among firms according to the cumulative distribution function $G_f(\cdot)$ that represents a Pareto distribution with scale parameter 1 and shape parameter θ_f .¹⁴ Technologies can be thought of as modes of firm organization or machines that consist of varieties of the intermediate good that require a specific mix of low- and high-skill labor to be operated. The Leontief structure of a particular production technology does not preclude firms from changing the relative amount of worker skill they employ in production. Rather, changing the relative skill employment requires employing a different technology.¹⁵

In order to serve the foreign market, a firm must pay f_x units of the final good as a per-period fixed cost of exporting. In addition, exporting entails a variable iceberg trading cost, $\tau > 1$. Thus, after paying f_x a firm can sell as much as it wishes of its variety in the Foreign market, but it must produce and ship τq_F units in order to sell q_F units abroad.

The total cost function for a firm with idiosyncratic productivity φ employing technology v , which will be referred to as a (φ, v) firm, that sells q_H and q_F units of its intermediate variety in the Home and Foreign countries is

$$C_{\varphi, v}(q_H, q_F; w) = \frac{q}{\varphi \mu_v(w)} + f_v + \mathbf{1}_{\{q_F > 0\}} f_x,$$

where $q = q_H + \tau q_F$ is the total quantity produced, $w = \{w_\ell, w_h\}$ are the wages of low- and high-skill

¹⁴As will become apparent in the technology calibration, the choice of scale parameter will not affect the results. The reason is that a change in the scale parameter will proportionally change all the technology-specific productivities.

¹⁵An alternative interpretation is that the firm chooses the quality of its intermediate variety rather than its production technology. In this interpretation, a firm increases the demand for its variety by increasing quality, so that α_v is a demand parameter rather than a technology parameter.

labor and $\mu_v(w)$ is the cost efficiency of technology v ,

$$\mu_v(w) = \frac{\alpha_v}{w_\ell + \lambda_v w_h}.$$

The cost efficiency of technology v times the firm's idiosyncratic productivity factor, φ , is the reciprocal of the unit cost of a (φ, v) firm.

Without loss of generality, the technologies are indexed in increasing order of fixed cost, so that if $v' > v$ then $f_{v'} > f_v$. In addition, $\alpha_0 = f_0 = 0$ so that a firm can always break even by choosing to produce nothing at zero cost. Because there are no sunk costs, a firm's decision in period t depends only on conditions prevailing in period t . In particular, firm optimization in period t depends only on current wages, $w_{\ell t}$ and w_{ht} , and current aggregate final good expenditure, R_t .¹⁶ The time subscript, t , is therefore suppressed in the firm-side analysis.

2.2.1 Firm Behavior

In each period, firms maximize profits by choosing prices, (p_H, p_F) , and associated quantities, (q_H, q_F) , in the Home and Foreign markets as well as a technology, v , from the menu of available technologies. In making the above choices, firms take as given aggregate expenditure on the final good, R , and wages of low- and high-skill labor, w_ℓ and w_h . The firm maximization problem can be solved in two stages. First, the firm solves for the optimal prices and quantities *given* technology choice, v . Second, it chooses the technology that yields the maximal profit.

Consider a firm with idiosyncratic productivity φ employing technology v . Given the CES structure of demand for intermediate varieties, the firm optimally sets prices, conditional on serving a given market, as a constant markup over marginal cost,

$$p_H(\varphi, v) = \frac{1}{\varphi \rho \mu_v(w)}, \quad p_F(\varphi, v) = \frac{\tau}{\varphi \rho \mu_v(w)}.$$

The optimal quantity *sold* in each market $i \in \{H, F\}$, conditional on serving market i , is

$$q_i(\varphi, v) = R p_i(\varphi, v)^{-\sigma}.$$

Variable profits in market i , conditional on serving this market, are then

$$\pi_i(\varphi, v) = A_i (\varphi \mu_v(w))^{\sigma-1},$$

¹⁶Because the equilibria are symmetric and the price of the final good is normalized to one, final good expenditure in the Home country is equal to final good production in the Home country.

where

$$A_H = \frac{R}{\sigma} \rho^{1-\sigma}, \quad A_F = \tau^{1-\sigma} A_H.$$

A firm will choose to export if and only if its variable profits in the Foreign country exceed the fixed exporting cost. Thus, the profit of a (φ, v) firm is

$$\pi(\varphi, v) = \pi_H(\varphi, v) + \max\{\pi_F(\varphi, v) - f_x, 0\}.$$

Any firm with idiosyncratic productivity φ will choose the technology that maximizes this profit,

$$v(\varphi) = \arg \max_v \pi(\varphi, v).$$

The more cost efficient a technology, the greater the savings it offers on variable cost. However, more cost-efficient technologies also require the payment of a higher fixed cost than less cost-efficient technologies.¹⁷ Therefore, the attractiveness to a firm of adopting a more cost-efficient technology at the expense of a higher fixed cost is increasing in the quantity produced by the firm. Since, for a given technology, the quantity produced by the firm is increasing in its idiosyncratic productivity, more productive firms adopt more cost-efficient technologies. Specifically, if $\varphi < \varphi'$ then $\mu_{v(\varphi)} \leq \mu_{v(\varphi')}$.

Denote by $\tilde{v} \in \{0, 1, \dots, \tilde{V}\}$ the subset of technologies from the full technology menu that are employed in the period under discussion, where this subset is indexed in increasing order of fixed cost, i.e. the order is unchanged from the full technology menu.¹⁸ The least productive active firms employ technology $\tilde{v} = 1$. As we move along the productivity space, firms will continue to use this technology until adopting technology $\tilde{v} = 2$ becomes more profitable. In general, technology \tilde{v} will be employed in the interval $\Phi_{\tilde{v}} = (\varphi_{\tilde{v}}, \varphi_{\tilde{v}+1})$, where firms with productivity $\varphi_{\tilde{v}}$ are indifferent between using technology \tilde{v} and $\tilde{v} + 1$ and $\varphi_{\tilde{v}+1} = \infty$. The productivity of the marginal technology upgrader, that is, the firm that is indifferent between technology $\tilde{v} - 1$ and \tilde{v} for some \tilde{v} , conditional on not switching export status, is

$$\varphi_{\tilde{v}} = \left[\frac{f_{\tilde{v}} - f_{\tilde{v}-1}}{A_H (1 + \mathbf{1}_{q_F > 0} \tau^{1-\sigma}) w_\ell^{1-\sigma} \left[(\alpha_{\tilde{v}} / (1 + \lambda_{\tilde{v}} \omega))^{\sigma-1} - (\alpha_{\tilde{v}-1} / (1 + \lambda_{\tilde{v}-1} \omega))^{\sigma-1} \right]} \right]^{1/(\sigma-1)},$$

where $\omega = w_h/w_\ell$ is the skill premium.

¹⁷In the following we ignore technologies, v , for which there exists a technology $v' < v$ such that $\mu_{v'} > \mu_v$. Technology v would never be adopted since it is dominated by technology v' which has both a lower marginal cost and a lower fixed cost. Therefore, with no loss of generality, we focus on the case that μ_v is increasing in v .

¹⁸In general, not all the available technologies will be used in every period. In Appendix A, I provide an algorithm for determining which technologies are used in equilibrium and by which firms.

The above equation shows the tradeoffs facing the firm in its technology choice. The numerator represents the added fixed cost involved in adopting the more cost-efficient technology, and the larger is this difference the more productive a firm needs to be to make it profitable to adopt the more cost-efficient technology. The opposite is true for the technology's Hicks-neutral component, $\alpha_{\tilde{v}}$.

Denote the market sizes of the Home and Foreign country facing the firm, normalized by the wage level, as $\tilde{A}_H = A_H w_\ell^{1-\sigma}$ and $\tilde{A}_F = A_H (\tau w_\ell)^{1-\sigma}$. The greater the normalized market size, the greater the quantity a firm will sell and therefore the more likely it is to adopt a more cost-efficient technology. Finally, if $\lambda_{\tilde{v}} > \lambda_{\tilde{v}-1}$, then a higher skill premium makes technology \tilde{v} less attractive relative to technology $\tilde{v} - 1$.

The tradeoff firms face in choosing their export status is similar to the tradeoff they face in choosing their technology. The benefit from gaining access to the Foreign market is increasing in firm productivity. However, since the cost of market access, f_x , is the same for all firms, it will be only the most productive firms that find it profitable to enter the export market. In particular, there exists a φ_x such that any firm with productivity $\varphi > \varphi_x$ exports and any firm with productivity $\varphi < \varphi_x$ serves only the Home market. The indifference of the marginal exporter between exporting and serving only the Home market implies, conditional on the marginal exporter not being also a technology upgrader,

$$\varphi_x = \left[\frac{f_x}{\tilde{A}_F (\alpha_{\tilde{v}_x} / (1 + \lambda_{\tilde{v}_x} \omega))^{\sigma-1}} \right]^{1/(\sigma-1)},$$

where \tilde{v}_x is the technology employed by the marginal exporter. If instead the marginal exporter is also a technology upgrader, the indifference condition yields

$$\varphi_x = \left[\frac{f_x + f_{\tilde{v}_x} - f_{\tilde{v}_x-1}}{\left[(\tilde{A}_H + \tilde{A}_F) (\alpha_{\tilde{v}_x} / (1 + \lambda_{\tilde{v}_x} \omega))^{\sigma-1} - \tilde{A}_H (\alpha_{\tilde{v}_x-1} / (1 + \lambda_{\tilde{v}_x-1} \omega))^{\sigma-1} \right]} \right]^{1/(\sigma-1)}.$$

The previous equations show that as the normalized Foreign market size, \tilde{A}_F , increases, the export cutoff decreases as it becomes more profitable to export.

These equations also illuminate how firm behavior is affected by trade liberalization. Anticipating the simulation results, trade liberalization, modeled as a decrease in τ , leads to a decrease in the normalized market size \tilde{A}_H , but an increase in both \tilde{A}_F and the sum $\tilde{A}_H + \tilde{A}_F$. For a given skill premium, the increase in \tilde{A}_F causes some firms that previously did not export to start doing

so. The increase in the world market size, $\tilde{A}_H + \tilde{A}_F$, leads all exporters to increase their output, and this induces some of them to upgrade their technology. However, firms that are not induced by trade liberalization to export face a smaller market as the influx of foreign varieties cuts into their market share. These firms, therefore, decrease their output and as a result some of them will downgrade their technology.

To the extent that exporters, on average, are more skill intensive than non-exporters, the increase in the market share of exporters relative to non-exporters will increase relative demand for high-skill labor. The increase in relative skill demand will, in turn, increase the skill premium. This general-equilibrium effect will make skill-intensive technologies relatively less attractive. Whether or not this will cause firms to upgrade or downgrade technology depends on the particular technology the firm employs, for, as the calibration will show, $\lambda_{\tilde{v}}$ is not monotonically related to \tilde{v} .

Since each firm chooses its export status and its technology, there are $2V$ possible configurations of the technology-export decision of a firm. The upshot of the preceding analysis is that firm optimization limits the possible configurations to at most $\tilde{V} + 1$. Firms employing technologies $\tilde{v} < \tilde{v}_x$ serve only the Home market, while firms employing technologies $\tilde{v} > \tilde{v}_x$ also serve the Foreign market. The set of firms using technology \tilde{v}_x can be partitioned into non-exporting firms, $(\varphi_{\tilde{v}_x}, \varphi_x)$, and exporting firms, $(\varphi_x, \varphi_{\tilde{v}_x+1})$. If $\varphi_x = \varphi_{\tilde{v}_x}$, that is, if the marginal exporter is also a technology upgrader, then there are \tilde{V} equilibrium technology-export configurations and otherwise there are $\tilde{V} + 1$.

2.2.2 Aggregation

The result that the firm productivity space can be divided into \tilde{V} intervals, one for each technology in use, makes it straightforward to aggregate firm-level factor demand and intermediate-good supply. The firm-level demand for labor is the per-unit labor requirement times the optimal quantity produced by the firm, derived in the previous section. The aggregate low- and high-skill labor demanded by firms are

$$L_f^d = \sum_{\tilde{v}} \int_{\Phi_{\tilde{v}}} \frac{(q_H + \mathbf{1}_{\varphi > \varphi_x} \tau q_F)}{\alpha_{\tilde{v}} \varphi} dG_f(\varphi), \quad H_f^d = \sum_{\tilde{v}} \int_{\Phi_{\tilde{v}}} \lambda_{\tilde{v}} \frac{(q_H + \mathbf{1}_{\varphi > \varphi_x} \tau q_F)}{\alpha_{\tilde{v}} \varphi} dG_f(\varphi),$$

where $\Phi_{\tilde{v}}$ is the interval of idiosyncratic productivities in which technology \tilde{v} maximizes profits.

The final-good demand by firms to cover the fixed technology cost and fixed cost of exporting is

$$F^d = \sum_{\tilde{v}} \int_{\Phi_{\tilde{v}}} (f_{\tilde{v}} + \mathbf{1}_{\varphi > \varphi_x} f_x) dG_f(\varphi).$$

The difference between variable profit and total fixed-cost expenditure yields aggregate firm profits,

$$\Pi = \sum_{\tilde{v}} \int_{\Phi_{\tilde{v}}} (\pi_H + \mathbf{1}_{\varphi > \varphi_x} \pi_F) dG_f(\varphi) - F^d.$$

Final-good supply is¹⁹

$$Q^s = \left[\sum_{\tilde{v}} \int_{\Phi_{\tilde{v}}} (q_H^\rho + \mathbf{1}_{\varphi > \varphi_x} q_F^\rho) dG_f(\varphi) \right]^{1/\rho}.$$

2.3 Workers

In this section, I develop an overlapping-generations model of endogenous skill acquisition. Workers are divided into age groups, $b \in \{1, 2, \dots, B + 1\}$. Age group $B + 1$ corresponds to death and the population of workers in age groups $b \leq B$, i.e. active workers, is unity. Workers are born uneducated into age group $b = 1$. In each period, a worker in age group $b \leq B$ dies with probability D_b , ages into age group $b + 1$ in the subsequent period with probability δ_b and remains in age group b in the subsequent period with probability $1 - \delta_b - D_b$. Upon birth each worker draws an idiosyncratic innate ability a from a Pareto distribution, $G_w(\cdot)$, with scale parameter 1 and shape parameter θ_w .²⁰ The measure of newborns is equal to the measure of deaths in each period so that there is no population growth.

In each period t , an uneducated worker faces a choice: work full time or pursue an education by becoming a student. If she works full time, the uneducated worker provides one unit of low-skill labor regardless of her innate ability and thereby earns $w_{\ell t}$. If instead she becomes a student, she works part time providing $m < 1$ units of low-skill labor and earns $mw_{\ell t}$. In addition, a student uses ℓ_e and h_e units of low- and high-skill labor, respectively, as teachers. The per-period cost of education,

¹⁹Because the countries are identical and the equilibria under consideration are symmetric, the quantity exported by Home firms to Foreign, q_F , equals the quantity imported by the Home country from the Foreign country. Therefore, the quantity following the indicator function indicates the quantity imported from the Foreign country.

²⁰As will become apparent, innate ability will only directly determine the income of high-earners. A Pareto distribution of abilities is therefore consistent with the empirical finding that the upper tail of the income distribution in the United States is well-approximated by a Pareto distribution. See, for example, Reed (2001).

then, consists of a tuition cost, $w_{\ell t} \ell_e + w_{ht} h_e$, and an opportunity cost of time, $(1 - m) w_{\ell t}$. A student in period t becomes educated in period $t + 1$ with probability δ_e and remains uneducated with the complementary probability.²¹

Once a worker becomes educated she remains educated for the remainder of her life. An educated worker with innate ability a provides a units of high-skill labor in each period and earns aw_{ht} in period t .²² Thus, there is a complementarity between innate ability and education, so that the benefit to education is increasing in innate ability.

Workers are hired both to produce intermediate goods for firms and as teachers to educate students. Workers can costlessly switch between teaching and working for an intermediate good firm, and between working for one firm or another. This ensures that in each period all workers of a given education level command the same wage per unit of labor. Workers are price takers, discount the future by $\beta < 1$, and their objective is to maximize their expected discounted lifetime consumption of the final good. Thus, in period s , workers' optimization decisions depend only on the current and future path of wages $\{w_{\ell t}, w_{ht}\}_{t=s}^{\infty}$ which, although endogenous, is exogenous from an individual worker's point of view.

2.3.1 Worker Optimization

Worker optimization boils down to a set of education decisions by uneducated workers. Given a path of wages, $\{w_{\ell t}, w_{ht}\}_{t=s}^{\infty}$, workers can compute the value of becoming a student and the value of working full time in any period.²³ Uneducated workers work full time if the latter exceeds the former and pursue an education if the reverse is true.

For workers of a given age group b , the higher the innate ability of the worker, the greater the benefits of acquiring an education. At the same time, the cost of acquiring an education is independent of the worker's innate ability. There therefore exists, for each period t and age group b , an education cutoff, a_{bt} , such that uneducated workers in age group b with ability $a > a_{bt}$ pursue

²¹The stochastic modeling of aging and education is to ensure tractability of the model by limiting the dimensionality of the workers' state space.

²²In principle, an educated worker can provide low-skill labor and thus must choose whether to provide one unit of low-skill labor or a units of high-skill labor. However, in equilibrium, educated workers always prefer to provide high-skill labor rather than low-skill labor, that is, $aw_{ht} > w_{\ell t}$ in all periods for all educated workers. The possibility that an educated worker chooses to provide low-skill labor is therefore ignored in the subsequent analysis.

²³In Appendix A I provide an algorithm for this computation.

an education and those with $a < a_{bt}$ work full time as uneducated workers. Thus, for any given path of wages, there is an associated path of education cutoffs, $\left\{ \{a_{bt}\}_{b=1}^B \right\}_{t=s}^{\infty}$, that fully characterizes workers' optimal policy.

In every period, $aw_{ht} > w_{lt}$ for all educated workers. Thus, the benefit to being an educated worker is increasing in the expected remaining lifetime of a worker, or, equivalently, decreasing in the worker's age group, b . However, the cost of education is independent of the worker's age. As a consequence, in every t the education cutoff is increasing in b , that is, if $b > b'$ then $a_{bt} > a_{b't}$.

The economy's adjustment to an increase in skill demand takes time for two reasons. First, skill supply takes time to adjust because education is a time-consuming activity and exactly how time consuming is governed by the parameter δ_e . Second, the adjustment takes time because, as just shown, older workers are less likely to pursue an education. Therefore, older workers, who would have pursued an education had present conditions prevailed in their youth, no longer find it profitable to do so. Only once these older workers die and are replaced by younger workers can the economy complete the adjustment to the increased skill demand.

2.3.2 Labor Supply

Every worker is characterized by a triplet (a, b, e) , where a is the worker's ability, b is the worker's age group and e is an indicator taking the value one if the worker is educated and zero otherwise. The worker distribution in period t is denoted by $W_t(a, b, e)$. An initial worker distribution, $W_s(a, b, e)$ and a path of education cutoffs, $\left\{ \{a_{bt}\}_{b=1}^B \right\}_{t=s}^{\infty}$, are sufficient to compute the full future path of worker distributions, $\{W_t(a, b, e)\}_{t=s}^{\infty}$.²⁴

The supply of labor and the demand for teachers in any period depend only on the worker distribution and education cutoffs in that period. Consider, therefore, any period with worker distribution $W(a, b, e)$ and worker policy $\{a_b\}_{b=1}^B$. The mass of students is the sum over all the age groups of uneducated workers with ability exceeding the education cutoff for that age group,

$$S = \sum_{b=1}^B \int_{a_b}^{\infty} W(a, b, 0) da.$$

Thus, the labor demand for teaching purposes is $S\ell_e$ and Sh_e for low- and high-skill labor, re-

²⁴In Appendix A I provide an algorithm for this computation.

spectively. The mass of full-time uneducated workers is calculated similarly except the sum is over workers with ability below the education cutoff for each age group,

$$L_{full}^s = \sum_{b=1}^B \int_1^{a_b} W(a, b, 0) da.$$

Therefore, the supply of low-skill labor is

$$L^s = L_{full}^s + mS$$

The supply of high-skill labor is the sum over all educated workers in each age group and each innate ability level weighted by this ability,

$$H^s = \sum_{b=1}^B \int_1^{\infty} aW(a, b, 1) da.$$

To summarize, given a path of wages $\{w_{\ell t}, w_{ht}\}_{t=s}^{\infty}$, worker optimization implies a path of worker policies in the form of education cutoffs for each age group, $\left\{\{a_{bt}\}_{b=1}^B\right\}_{t=s}^{\infty}$. This path of worker policies, together with an initial distribution of workers, yields the entire future path of worker distributions, $\{W_t(a, b, e)\}_{t=s}^{\infty}$. Finally, the worker distribution and education cutoffs in any period yield the low- and high-skill labor supply as well as the demand for teachers in that particular period.

2.4 Equilibrium

An equilibrium of the open economy for an initial worker distribution $W_s(a, b, e)$ is characterized by paths of wages, $\{w_{\ell}, w_h\}_{t=s}^{\infty}$, and aggregate final good expenditure, $\{R\}_{t=s}^{\infty}$. These paths are sufficient to compute optimal worker and firm policy for every period. The paths of wages and final-good expenditure are an equilibrium if, given the optimal worker and firm policies implied by these paths, the labor market and goods market clear in every period.

Labor-market clearing requires that the spot market for each type of labor clears, that is, the sum of labor demand by students and firms for each skill equals its supply,

$$L^s = S\ell_e + L_f^d, \quad H^s = Sh_e + H_f^d.$$

Product market clearing requires that the quantity supplied of the final good equals expenditure on the final good. Expenditure on the final good is the sum of firm demand for the final good used to

cover fixed costs and worker expenditure on the final good, which is, in turn, the sum of aggregate profits and worker wage income net of tuition costs,

$$Q = R = F^d + \Pi + w_\ell L^s + w_h H^s - S(w_\ell \ell_e + w_h h_e).$$

A steady-state equilibrium is one in which all of the endogenous variables are constant over time.

3 Calibration

The goal of this paper is to provide estimates of the effects of counterfactual trade policies. With the theoretical model in hand, the next stage, therefore, is to calibrate the model to data. In this section, I describe how I calibrate the parameters of the theoretical model to match both macro and micro features of the US economy in 2007.

3.1 Data

Worker data comes from the Integrated Public Use Microdata Series - Current Population Survey (IPUMS-CPS), a publicly available dataset that consists of a random sample of the March supplement of the Current Population Survey (CPS). The unit of observation in this dataset is the individual. The CPS provides information on a wide array of individual level data. In particular, I use data on age, income, educational attainment and labor market participation.

The CPS provides data on students only if they are under the age of twenty-four. Therefore, for information on students, I supplement the CPS data with information from the IPUMS - American Community Survey (ACS) sample. Department of Education data provides information on average time to degree completion and average tuition cost of a post-secondary degree.

The CPS data also provide information on the industry and size of the firm in which the individual works. Firm size in the data is divided into six bins. The smallest size bin consists of firms with one to nine workers and the largest bin is for firms with 1000 or more workers. Because the CPS data also provides the educational attainment of the individual, this data is informative about the relative demand for educated workers by firm size and industry.

The remainder of data on firms comes from the Economic Census and the COMPUSTAT database. The Economic Census divides the size distribution of firms into twenty-five bins and provides information on the number of firms and total payroll of firms in each bin. The smallest bin is of firms with one to four workers, while the largest bin is of firms with 10,000 or more workers. Each of these bins, except the largest bin, will be used to calibrate a single technology from the technology menu in the theoretical model. Since the bin of the largest firms represents approximately 40% of all employment and because this represents the largest bin in terms of size range, it is important to allow for more than one technology in this bin. I therefore supplement the Economic Census data with COMPUSTAT data which provides information on firm size for all publicly-traded firms in the United States. I use this data to add eight bins for those firms that employ more than 10,000 workers, which allows me to calibrate eight technologies for these firms. To avoid confusion, I will refer to the Economic Census and COMPUSTAT bins as size bins or simply bins, while I will refer to the bins from the CPS data as CPS bins.²⁵

In the baseline calibration, I interpret educated workers in the model to be workers who have completed at least a BA degree. Individuals with no post-secondary degree correspond to uneducated workers in the model. Individuals with a degree from a two-year post-secondary institution, such as an Associate degree, are weighted so that they count as half an educated worker and half an uneducated worker.²⁶

In my data analysis, I consider only workers between the ages of nineteen and sixty-eight. Nineteen roughly corresponds to the first year of college for most college attendees, and by age sixty-eight only a small percentage of the population is still in the labor force. The intermediate-good sector is interpreted as the universe of all active firms in the United States in 2007 in all industries except government services and education services. The education services sector is included in the analysis to calibrate the labor requirement for teaching.

²⁵See Appendix B for details on the size bins and CPS bins.

²⁶Alternative interpretations of who is educated and who is uneducated do not appear to alter the results in any meaningful way. For example, excluding two-year post-secondary degree recipients from the pool of educated workers; using alternative weights for these workers; or including, in the educated worker pool, workers who have completed some college does not affect the results in this paper.

3.2 Calibration Details

In the baseline simulations each period corresponds to one year. Each age group corresponds to a five-year window (19-23, 24-28, etc.) so that in total there are ten age groups ($B = 10$). The probability of aging, δ_b , for any age group is taken to be 0.2. This ensures that on average workers spend five years in each age group. Similarly, the probability of a student becoming educated, δ_e , is taken to be 0.22 to ensure that on average it takes four and a half years to complete a BA degree as per Department of Education data. The probability of death in each age group, D_b , is chosen to match the decrease in labor market participation as workers age. Finally, the discount factor β is taken to equal 0.97 to match a real annual interest rate of 3%.

I take the elasticity of substitution between varieties of the intermediate good to be 3.8. This matches the estimate in Bernard, Jensen, Eaton and Kortum (2003). The theoretical model implies that the share of exports in total shipments of a firm, conditional on exporting, is $\tau^{1-\sigma} / (1 + \tau^{1-\sigma})$. This expression is set equal to 0.14 to match the fact, reported in Bernard, Jensen, Redding and Schott (2007), that on average exporters ship abroad 14% of their total shipments. The resulting iceberg cost is $\tau = 1.91$.²⁷

The remaining parameters that require calibration are (1) the firm level technological parameters, $\{\alpha_v, \lambda_v, f_v\}_{v=1}^V$ and f_x ; (2) the teaching requirements, ℓ_e and h_e ; (3) the mass of firms, M ; and (4) the Pareto shape parameters for worker innate ability, θ_w , and firm idiosyncratic productivity, θ_f .

3.2.1 More on Calibration

I calibrate the Pareto shape parameters, (θ_w, θ_f) , by simulating a steady-state equilibrium in the model and matching some of the simulated equilibrium moments to macro moments in the data. The rest of the parameters are calibrated in the process as well, as described below. For a given pair, (θ_w, θ_f) , I solve the model for a steady-state equilibrium by iterating over values of wages and aggregate final-good expenditure, (w_ℓ, w_h, R) .²⁸ I choose the Pareto parameters so that the relative skill supply and the average supply of high-skill labor per educated worker predicted by the model

²⁷This is in line with, for example, the iceberg trade cost used in Melitz and Redding (2013) of 1.83. In general, this level of trade costs is consistent with the trade-cost estimates in the literature. For a review of these estimates, see Anderson and van Wincoop (2003).

²⁸The low-skill wage is given by data because all uneducated workers provide one unit of low-skill labor. Thus, w_ℓ is equal to the average wage of uneducated workers.

match the data.

A low Pareto shape parameter is associated with a fat right tail of the distribution. The lower the value of θ_w , the more units of high-skill labor each individual educated worker provides, on average. Therefore, lower values of θ_w imply more units of high-skill labor for a given number of educated workers. This fact allows me to identify θ_w from the relationship between the number of educated workers and the supply of high-skill labor, i.e. the average supply of high-skill labor per educated worker. I delay the discussion of θ_f until after the description of the firm-level technology calibration at which point its role in the calibration will become clearer.

I describe here the remainder of the calibration for given values of $\theta_f, \theta_w, w_\ell, w_h$ and R . The labor requirement for teaching is calibrated to match a yearly tuition cost of 20,000 dollars as per Department of Education data and relative expenditure on low- to high-skill labor in the education sector of 0.28 as calculated from CPS data,

$$\ell_e w_\ell + h_e w_h = 20,000, \quad \frac{\ell_e w_\ell}{h_e w_h} = 0.28.$$

Given ℓ_e and h_e , I compute the skill supply to the intermediate-good producing firms, that is, the skill supply net of demand for teachers. I then calibrate the firm-level technological parameters with a procedure I describe in the next section. This yields the relative demand for skills by firms. I set the mass of firms, M , so that the *level* of firms' skill demand equals the *level* of supply of skill to firms.

3.2.2 Calibration of Firm-Level Technologies

In this section I briefly describe the calibration of firm-level technologies. The full details of the calibration can be found in Appendix B.

In order to calibrate the firm-level technologies, I assume that each of the thirty-two size bins described in Section 3.1 represents a unique technology, so that all firms in a bin adopt the same technology and $V = 32$. The implication of this assumption is that the smallest firm in a bin, say bin v , represents the marginal firm that is indifferent between using technology $v - 1$ and upgrading to technology v . Denote the size of this firm s_v .

For expositional clarity, suppose that the CPS bins and Economic Census bins were identical.

The CPS data provides information on relative expenditure on labor in each bin which, given wages, implies a skill intensity, λ_v , for each technology. This choice of λ_v ensures that the model matches the data on relative expenditure by education level within each CPS bin. In addition, the assumed functional form of the firm-productivity distribution together with data on the proportion of firms in each bin imply a productivity, φ_v , for firms on the bin boundaries.

Given a firm's idiosyncratic productivity and its choice of technology, firm optimization implies a unique size for the firm. In the calibration, the size of the firm on the bin boundary is known, s_v , and by assumption this firm employs technology v . The parameters φ_v and λ_v are also known as described above. Therefore, inverting the optimization problem of the firm yields the technology-specific productivity α_v .

The only remaining parameters are the fixed costs. Consider a firm of size s_1 , the smallest active firm. This firm must earn zero profits as it must be indifferent between producing and being inactive. The fixed cost, f_1 , must, in order to ensure zero profits, equal the firm's variable profits, which can be computed from previously calibrated parameters. Now, suppose the fixed costs have been calibrated for all technologies up to $v - 1$. The firm of size s_v must be indifferent between using technology $v - 1$ and technology v . The fixed cost of using technology v , f_v , must, therefore, equal the variable profit from employing technology v less the total profit from using technology $v - 1$, a known number since f_{v-1} has already been calibrated.

I assume that the marginal exporter, that is, the smallest exporting firm, is also a technology upgrader.²⁹ The productivity of the marginal exporter, φ_x , is chosen so that roughly 4% of firms export as reported by Bernard, Jensen, Redding and Schott (2007). This, however, presents a complication as the procedure described above would then yield only the sum of the fixed cost of the technology adopted by the marginal exporter and the fixed cost of exporting.

Because the marginal exporter both upgrades technology and exports, and because of the complementarity between more productive technologies and exporting, it must be the case that absent the technology upgrade the firm would prefer not to export. This provides a lower bound for f_x as it must be sufficiently large to ensure that such a firm would indeed not export. Conversely,

²⁹The assumption that the marginal exporter also upgrades technology is in line with evidence that firms induced by trade liberalization to export also upgrade technology. See, for example, Lileeva and Treffer (2010) and Bustos (2007).

the firm must prefer not to upgrade technology if it does not export. This implies a lower bound for the fixed technology-usage cost. The sum of the fixed costs together with these lower bounds imply a range for each of these fixed cost. The fixed cost of exporting, in 2007 US dollars, is in the interval from 98,425 to 108,935. In the baseline parametrization I take f_x to be the midpoint of this interval.³⁰

3.2.3 Results of Technology Calibration

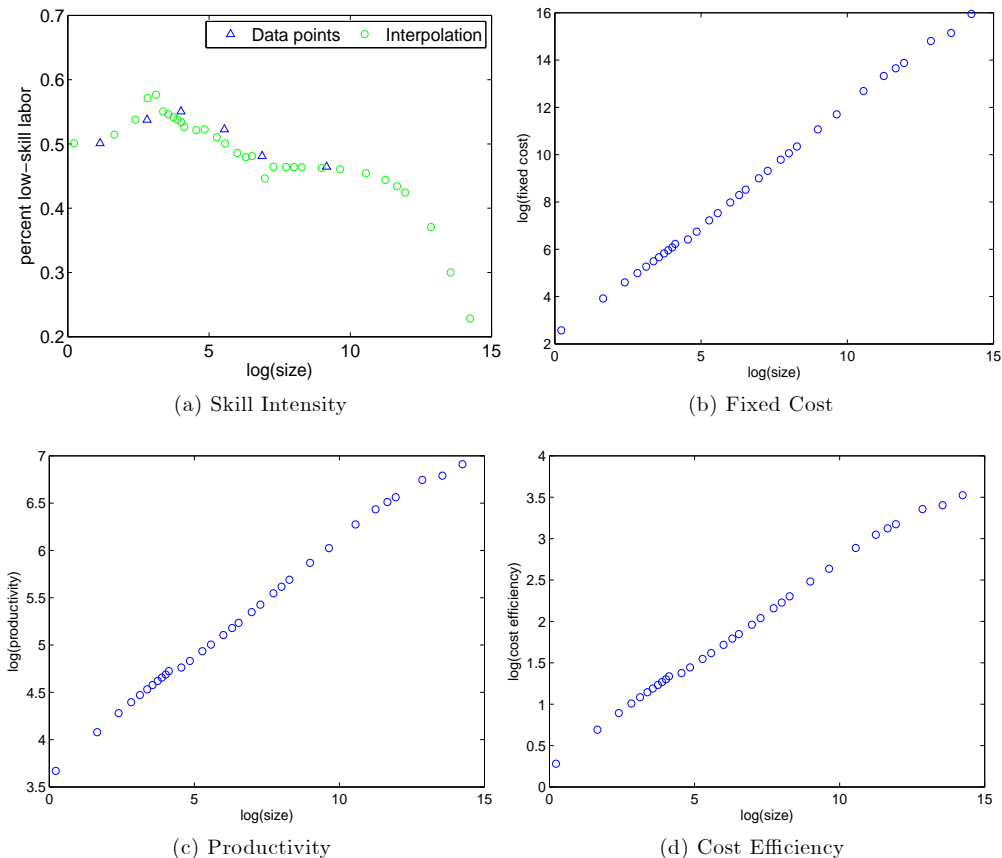
The results of the technology calibration are summarized in Figure 1. One of the advantages of the technology calibration is that it does not require any assumptions about the relationship between $(\alpha_v, \lambda_v, f_v)$ across technologies. Other studies of the effect of trade on relative skill demand have generally assumed a positive and monotonic relationship between firm productivity and skill intensity. Burstein and Vogel (2011) study a model with no technology choice but assume that a firm’s skill intensity is increasing in its idiosyncratic productivity. Bustos (2011) studies a model in which firms choose between two technologies, the more productive of which is more skill intensive. While this is a reasonable assumption when there are only two technologies, as more productive firms are, *on average*, more skill intensive than less productive firms, with more technologies this assumption need not hold. Indeed, the result of the calibration is not a monotonic relationship between firm productivity and skill intensity. Instead, the data suggest, as seen in Figure 1a, that there exists an approximately inverted-U relationship between firm productivity and the skill intensity of the technology it employs.

It is apparent from Figure 1a that the inverted-U relationship between firm productivity and skill intensity is a feature of the data and not of the specific interpolation method I use.³¹ Indeed, the data points in the figure are unimodal, and it is only the interpolation that generates the slight deviations from unimodality. These deviations are a result of the fact that the interpolation method I employ does not impose monotonicity of the relationship between firm productivity and skill intensity between data points. In order to assess the robustness of my results to alternative interpolation techniques, I fit fifth and sixth degree polynomials to the data points in Figure 1a

³⁰Other choices of f_x in this interval do not influence the results.

³¹In reality, the Economic Census bins are a non-trivial partition of the CPS bins. In order to identify the skill intensity of each technology within the CPS bins, I use an interpolation technique detailed in Appendix B.

Figure 1: Menu of Technologies



Notes: Panel (a) in this figure plots low-skill labor used as a percentage of total units of labor used in each technology. The data points are calculated directly from CPS data and the interpolation procedure is described in Appendix B.

instead of the interpolation used in the baseline calibration described in Appendix B. In both cases, the alternative interpolations have little quantitative impact on the results I report in Section 4.

The theoretical model implies that, among the technologies observed in equilibrium, the cost efficiencies, $\mu_v = \alpha_v / (w_\ell + \lambda_v w_h)$, are increasing in v , i.e. $\mu_v < \mu_{v'}$ for $v < v'$. The reason is that more-productive firms benefit more than less-productive firms from more cost-efficient technologies. Therefore, if a given firm prefers technology v to v' , then a more productive firm would as well, and technology v' would not be observed in equilibrium. For similar reasons, the theoretical model precludes the possibility that $f_v > f_{v'}$ for some $v' > v$ used in equilibrium.

The calibration, however, precludes neither the possibility that $f_v > f_{v'}$ nor the possibility that $\mu_v > \mu_{v'}$ for some $v' > v$. Nevertheless, despite the fact that it is not assumed, the calibration does

deliver increasing fixed cost and increasing cost efficiency in v , as is evident from Figures 1b and 1d. Thus, the model, and in particular the firm side of the model, is consistent with the dimensions of the data used in the calibration.

3.3 Discussion of Calibration

After calibration of the firm-level technologies, we are in a position to understand the role of the Pareto shape parameter, θ_f , that governs the distribution of firm idiosyncratic productivities. The calibration of the firm-level technologies ensures that the model matches the data on relative demand for skills within each CPS bin. However, this does not guarantee that the aggregate relative demand for skills will match the data. The reason is that the technology calibration treats each CPS bin separately and as such does not attempt to match any relationship across CPS bins. Calibration of θ_f to match the aggregate relative demand ensures that each CPS bin represents roughly the same proportion of total employment in the model as in the data.

The calibration yields $\theta_f = 9.3$, a value that is higher than the one used in some other studies. For example, Melitz and Ghironi (2005) use a value of 3.4 to match a standard deviation of log of firm sales of 1.67 using the same elasticity of substitution between varieties as I do, $\sigma = 3.8$. In their model, the idiosyncratic productivity is the only source of size variability across firms. A higher value of θ_f would imply that there is less size variability if the idiosyncratic productivity were the only source of size differences among firms. However, in the current model, the differences in size between firms resulting from differences in idiosyncratic productivities are amplified by firms' technology choices. The amplification results from firms with higher idiosyncratic productivities adopting more-productive technologies. Thus, differences in idiosyncratic productivities need only account for part of the overall size variability in the data. In order to match the empirical size distribution, the model does require a higher θ_f than a model in which idiosyncratic productivity is the only source of size difference across firms.

Since I calibrate the technologies at the firm-level to a 2007 cross section, a potential concern is whether the calibration matches the aggregate elasticity of substitution between high- and low-skill labor. Katz and Murphy (1992) estimate that in the United States a 10% increase in relative skill

supply leads to a 7% decrease in the skill premium, a ratio of 1.4 with a standard error of 0.3. This is the elasticity of substitution between high- and low-skill labor if, as is assumed in that paper, the aggregate production function has a constant elasticity of substitution. In my model, however, the aggregate production function does not exhibit a constant elasticity of substitution. Nevertheless, in order to assess the performance of the calibrated model in this dimension, I mimic the analogous statistic in the calibrated model by computing the change in the equilibrium skill premium in response to an exogenous shock to the relative skill supply. I find that a 10% change in relative skill supply leads to a 9% change in the skill premium, a ratio of 1.1. This is well within the confidence interval estimated by Katz and Murphy (1992).

To understand why the ratio I compute is lower than the point estimate of 1.4 in Katz and Murphy (1992), note that they make the identification assumption that relative skill supplies do not respond to shocks to relative skill demand. However, if skill supplies are endogenous, as in my model, then this assumption is unlikely to hold, and an observed increase in relative skill supply is likely to be, at least in part, a response to an increase in skill demand. Absent the increase in skill demand, the same increase in skill supply would imply a larger decrease in the skill premium than if there were an increase in skill demand. Thus, computing the aforementioned ratio from data in which the increase in relative skill supply is accompanied by an increase in relative skill demand, as in Katz and Murphy (1992), leads to a higher estimate of this ratio than when the change in relative skill supply is not accompanied by an increase in relative skill demand as in my computation.

4 Quantitative Analysis

With the theoretical model now calibrated, the stage is set for conducting the quantitative analysis. My goal is to study the effects of a counterfactual bilateral removal of policy barriers to trade, modeled as a reduction in the iceberg cost, τ . Of course, trade policy by itself cannot remove all iceberg costs as some of this cost is a real resource cost determined by the available transportation technology and the distance between countries which policy cannot eliminate. The policy I consider, therefore, is the removal of border barriers to trade, i.e. costs that are not attributable to distance. Anderson and van Wincoop (2003) find that the border barrier between Canada and the US is 26%.

In line with this finding, I simulate the economy's response to an unanticipated once-and-for-all removal of border barriers equal to 26% which implies a reduction of τ from 1.91 to $1.91/1.26 = 1.51$.

In order to solve for the entire equilibrium transition path, I employ a shooting algorithm. The algorithm involves solving first for the initial and terminal steady states and then iterating on a path of wages until convergence. For a given guess of the wage path, worker values are calculated by iterating backwards from worker values in the terminal steady state. These worker values together with the initial worker distribution allow me to iterate forward to compute the worker distribution and skill supply in all periods. I then compute the implied equilibrium path of wages given this skill supply, which will, in general, differ from the initial guess of the path of wages. I then adjust the guess for the path of wages and repeat the above steps. The algorithm converges when the guess of the path of wages coincides with the computed equilibrium path of wages given the skill supply implied by the guess. The algorithm is laid out in detail in Appendix C.

4.1 Results

4.1.1 Increase in Relative Skill Demand

The catalyst for the adjustment in skill supply, the source of all the dynamics in the model, is the change in relative skill demand precipitated by the new trade policy.³² I therefore begin the analysis with a brief description of the effect of trade liberalization on relative skill demand. In the calibration, I do not impose assumptions on the technology menu that guarantee that trade liberalization induces an increase in relative skill demand. Nevertheless, relative skill demand does increase following trade liberalization. In fact, if skill supplies were counterfactually fixed at their pre-liberalization level, the skill premium would increase by approximately 12%.

The increase in the relative skill demand is a result of the forces discussed in Section 2. Exporters perceive trade liberalization as an increase in market size while non-exporters perceive it as a decrease in market size. This leads to a reallocation of production shares from non-exporters to exporters. Since, on average, exporters employ more skill-intensive technologies than non-exporters, this reallocation leads to an increase in relative skill demand.³³

³²If skill supplies were fixed, there would be no dynamics and the economy would immediately reach its post-liberalization steady state.

³³This is not a general feature of the model. In particular, because of the inverted-U relationship between firm productivity and skill intensity, it is not necessarily the case that exporters are on average more skill-intensive than

Firms' choices of technologies following trade liberalization are determined by two forces. First, the increase in market size for exporters induces some of them to upgrade technology for a given skill premium. Second, trade-induced changes in the skill premium lead to changes in the relative cost efficiency of the technologies from which firms choose.

For concreteness, consider the measure of firms employing the most productive technology. In the first year after trade liberalization, the measure of these firms decreases by 20%, while in the long run there is an 11% increase relative to pre-liberalization. Initially, because of the policy-induced increase in the skill premium, the skill-premium effect dominates the market-size effect. However, in the long run, the market-size effect dominates as increasing skill supplies depress the skill premium.³⁴

4.1.2 Income Paths

My main concern is with the division of the gains from trade among workers and the inequality caused by the new trade policy. However, it is useful to begin the analysis with a description of income paths, since they ultimately are the determinants of the gains from trade and inequality. Figure 2 charts the evolution of average income in the economy from the period immediately preceding the removal of trade barriers until forty years after their removal.³⁵

If this were the whole story, we might reasonably conclude that although trade has important consequences, the dynamics appear to be of minor importance. Indeed, average income in the economy increases by 5.4% in the first year after trade liberalization, but this already represents 87.5% of the long-run increase in average income. However, the average income conceals the differential impact of trade on workers of different education levels.


Figures 3a and 3b show the decomposition of the average income in the economy into the average income among uneducated and educated workers. These graphs reveal that the relatively muted dynamics observed in the aggregate average income results from the fact that educated and

non-exporters. In general, this depends on the precise location in the firm distribution of the export cutoff. If the export cutoff is in the increasing part of the inverted-U, then trade liberalization may actually lower the relative demand for skill.

³⁴Changes in the skill premium incentivize workers to pursue an education which increases the demand for teaching labor. Because the labor teaching requirement is more skill intensive than the aggregate skill intensity in the intermediate-good sector, increases in the number of students increase the relative demand for skill. Although this effect exists in the model, it is not quantitatively important and I therefore focus on the shifts in relative skill demand within the intermediate-good sector.

³⁵Because the price of the final good is normalized to unity in every period, all wages are real.

Figure 2: Percentage Change in Average Income Relative to Pre-liberalization Along the Transition Path



Notes: The path shown is the average income over all full-time workers in the economy from the first year before trade liberalization until forty years into the transition.


uneducated workers' average incomes evolve in opposite directions after the initial trade shock. The increase in relative demand for high-skill labor creates a shortage of educated workers that sharply drives up their wages to beyond the long-run steady-state level. The average income of educated workers then falls back towards its long-run level as workers respond to the high wages by pursuing an education and thus augmenting the supply of high-skill labor.

For uneducated workers the situation is reversed. The increase in relative demand for high-skill labor means that there is a surplus of uneducated workers relative to the long run. Their wage, therefore, increases gradually together with the decrease in relative supply of low-skill labor. Despite the surplus of low-skill labor, uneducated workers' income does not decrease initially because trade liberalization leads to an increase in the varieties of the intermediate good available to all workers, which increases their real income.

The lack of significant dynamics in the aggregate thus masks the fact that the fortunes of educated and uneducated workers are moving in opposite directions during the transition. The differential impact on the two types of workers roughly offset one another in the aggregate. Indeed, this will be a recurring theme in what follows.³⁶

³⁶ A close inspection of the paths of the average incomes of educated and uneducated workers makes it apparent that the average incomes of both high- and low-skill workers oscillate around the long-run steady-state average incomes

Figure 3: Percentage Change in Average Income Relative to Pre-liberalization by Education Along the Transition Path



Notes: The figures chart the evolution of average incomes among full-time workers by education from one year before trade liberalization until forty years into the transition.

4.1.3 Gains From Trade

We are now ready to turn to the main results of the paper, beginning with an analysis of the gains generated by trade liberalization. These gains are defined as the percentage increase, compared to pre-liberalization, in expected discounted real income net of education costs for present and future generations of workers. The simulation results show that the removal of policy trade barriers leads to aggregate gains from trade of 5.9%. However, these gains from trade are not equally shared by all workers. The gains from trade are most unequally divided among workers who are alive at the time of trade liberalization, while for subsequent generations the gains are divided more equitably.

Workers Alive at the Time of Trade Liberalization

Table 1 reports the gains from trade for workers by education and across age groups for workers who are alive at the time of trade liberalization. The last row shows that although the youngest workers as a group gain the least from trade liberalization, the gains are shared quite evenly between age groups when gains are averaged over all education groups. However, the last column shows that the gains are more uneven along the worker education dimension. Educated workers as a group gain

until convergence. These cycles are a direct result of the alternating periods of scarcity and surplus of high-skill labor, which I will return to in the discussion of the skill-supply adjustment.

the most, with an increase of 7.7% in their expected discounted lifetime earnings, while uneducated workers who are not induced to acquire an education gain the least at 4.1%.

Table 1: Gains from Trade for Workers Alive at Implementation of Trade Liberalization

Group	19-23	34-38	49-53	64-68	Total
Uneducated	4.6	4.3	3.5	1.5	4.1
Educated	7.2	7.5	8.2	9.9	7.7
Students	6.5	6.6	6.8	7.7	6.6
Switchers	5.1	5.0	4.6	4.0	5.1
Total	5.8	6.0	6.0	6.0	6.0

Notes: The columns represent age groups. Switchers refers to uneducated workers who were induced by trade liberalization to pursue an education. Uneducated workers refers to workers who will never pursue an education.

Although age groups as a whole gain relatively evenly from trade, this masks the differences between education levels within age groups and within education levels across age groups. For example, the older an educated worker, the greater her gains from trade. The situation is reversed for uneducated workers who gain less the older they are. The older a worker at the time of trade liberalization, the larger the proportion of their remaining lifetime is spent with a relatively high educated-worker wage and low uneducated-worker wage. This works in favor of older educated workers, but is particularly detrimental to old uneducated workers.

Thus, among age groups, the gains are most unevenly divided among the oldest workers. Educated old workers gain 9.9%, while uneducated old workers gain only 1.5%. For the youngest age group the inequity in the division of the gains from trade is smaller but still pronounced. Uneducated workers gain 4.6%, while educated workers gain 7.2%.

The upshot of these results is that there is considerable heterogeneity in how workers alive at the time of implementation of the new trade policy are affected by the policy. An individual worker's gains depend on both her education and age. In general, educated workers gain more than uneducated workers. Among educated workers the gains are larger for older workers, while the reverse is the case for uneducated workers.

Worker Born After Trade Liberalization

The ex-ante gains from trade, that is, the expected gains at the time of birth, are much more evenly divided among workers born after trade liberalization. There are two reasons for this. The first is that workers born after trade liberalization are much less affected by the large initial rise in the returns to education. The second reason is that all workers are born uneducated. Much of the inequity in the division of the gains from trade is caused by the fact that workers who had already acquired an education pre-liberalization are rewarded with a particularly large windfall.

Table 2: Gains for Workers Born After Trade Liberalization by Birth Cohort

Group	Impact	5 years	10 years	25 years	Steady State
Uneducated	4.6	5.2	5.5	5.7	5.6
Students	6.5	6.3	6.2	6.2	6.3
Total	5.5	5.7	5.8	5.9	5.8

Notes: Uneducated workers refers to workers who will never pursue an education. Students refers to workers who begin their life as students, although they may stop pursuing an education if unsuccessful at acquiring an education.

Table 2 reports the gains from trade, by birth cohort, for workers born after the new trade policy is implemented. The percentages reported in the table are a worker's ex-ante gains from trade at birth taking into account that even if they choose to pursue an education, they may not be successful in the endeavor. Of course, the gains among students will be divided more unequally ex-post as some students will become educated while others will not.

The last row in Table 2 shows that the gains from trade averaged over all workers in a cohort are increasing with time, although the increase is minor. Indeed, the small increase reflects the fact that most of the aggregate gains are realized immediately, while the important dynamics are reflected in the division of these gains across education groups.

The first column of the table shows that workers whose innate ability is sufficiently low so they will never pursue an education gain 4.6% relative to their discounted lifetime earnings had the new trade policy not been implemented. In contrast, the average gains from trade for workers who do pursue an education is 6.5%. The difference in gains for the two groups diminishes along the transition path. Nevertheless, even in the post-liberalization steady state the gains for workers who pursue an education are higher than for those that do not, although the difference is only 0.7 percentage points.

4.1.4 Earnings Inequality

In this section, I consider the cross-sectional differentials in lifetime earnings among workers along the transition path. This is distinct from the differences in the gains from trade considered above because in this section I take into account the ex-post outcomes of workers' education decisions and not just their ex-ante expected discounted earnings. In addition, I consider workers in all age groups and not just the gains for workers at the time of their birth. For brevity, I will refer to the net present value of a worker's lifetime income stream net of education costs simply as lifetime earnings. In addition, when I discuss inequality between groups of workers I will be referring to relative average lifetime earnings between these groups.³⁷ Thus, for example, the college premium is defined here as the average discounted lifetime earnings of an educated worker relative to the average discounted lifetime earnings of an uneducated worker that will never pursue an education.

Inequality By Education

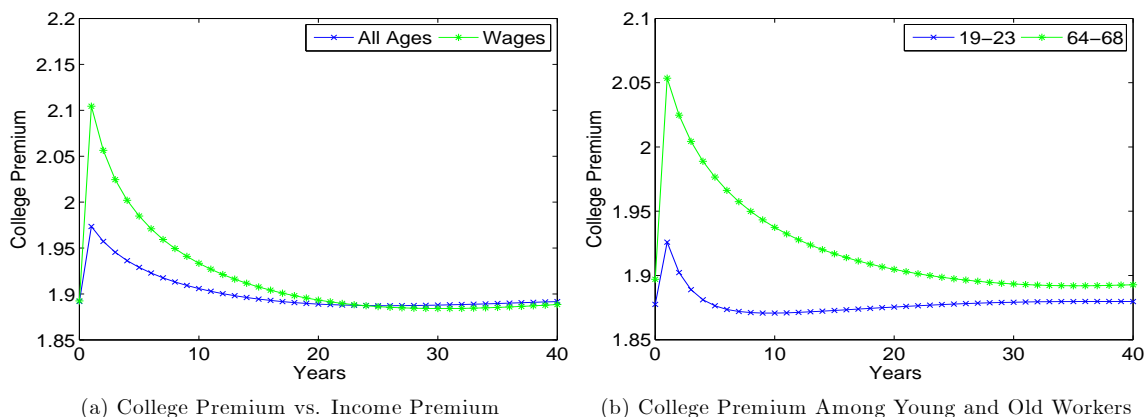
Figure 4 plots the evolution of the college premium from the steady state prevailing before trade liberalization until forty years into the transition. For all age groups there is a large immediate increase in earnings inequality between educated workers and uneducated workers followed by a gradual decline.

For the youngest workers, aged 19-23, lifetime-earnings inequality, measured here by the college premium, actually falls slightly below its pre-liberalization level approximately five years into the transition. There are two reasons for this. First, the influx of educated workers causes the relative returns of education to fall. Second, because of the increased incentives to pursue an education, the marginal worker is of lower ability after liberalization than before, which lowers the average ability of educated workers. Therefore, for a given wage per unit of high-skill labor, the average earnings per educated worker drops.

Next, consider the difference in the college premium for young and old workers. This difference is the sum of two effects. First, as workers age, only the most able continue to pursue an education while the least able discontinue their educational pursuit, which ensures that the average ability of

³⁷Care should be taken in the interpretation of this inequality in terms of welfare inequality as it does not account for the fact that educated workers have paid for their education in the past.

Figure 4: College Premium Along the Transition Path



Notes: The college premium is defined as the lifetime earnings of an educated worker relative to that of an uneducated worker. The income premium is defined as the within period income of an educated worker relative to an uneducated worker. Panel (a) plots the college premium across all age groups as well as the income premium. Panel (b) plots the college premium within the youngest and oldest age groups.

workers is increasing with age, for a given skill premium. Thus, even in the steady state, earnings inequality is highest among the oldest workers. This effect tends to increase the difference in the within-age-group college premium between young and old.

Second, older workers' lifetime earnings depend only on wages in the near future whereas younger workers' lifetime earnings depend also on wages in the more distant future. This will increase the difference in the college premium between old and young only if the skill premium is decreasing. Immediately after the transition this leads to an increase in the difference in inequality between age groups, but after the wages overshoot their long-run level, as seen in Figure 3, the effect is reversed. Nevertheless, Figure 4b shows that earnings inequality is always higher among the older workers.

Table 3 gives a sense of the magnitude of the changes in inequality following trade liberalization. The lifetime earnings of old educated workers relative to old uneducated workers are 8.2% higher immediately after trade liberalization compared to before trade liberalization, while for young workers the increase is smaller at 2.6%. In the long run, the endogenous skill-supply response brought about by the high returns to education depresses the college premium. Hence, in the new steady state, inequality is only slightly higher than in the pre-liberalization steady state.

It is useful to contrast the change in the long-run college premium with the results of Findlay and Kierzkowski (1983). Findlay and Kierzkowski (1983) study the long-run effects of introducing

Table 3: Percentage Change in the College Premium After Trade Liberalization

	Impact	5 years	10 years	25 years	Steady State
College Premium, All Ages	4.3	2.0	0.7	-0.2	0.1
College Premium, 19-23	2.6	-0.1	-0.4	0.0	0.1
College Premium, 64-68	8.2	4.2	2.1	0.0	0.1
Income Premium	11.2	4.9	2.2	-0.3	0.1

Notes: The college premium is defined as the average discounted lifetime earnings of educated workers relative to that of uneducated workers who will never pursue an education. The income premium is the average income of an educated worker relative to an uneducated worker within a period.

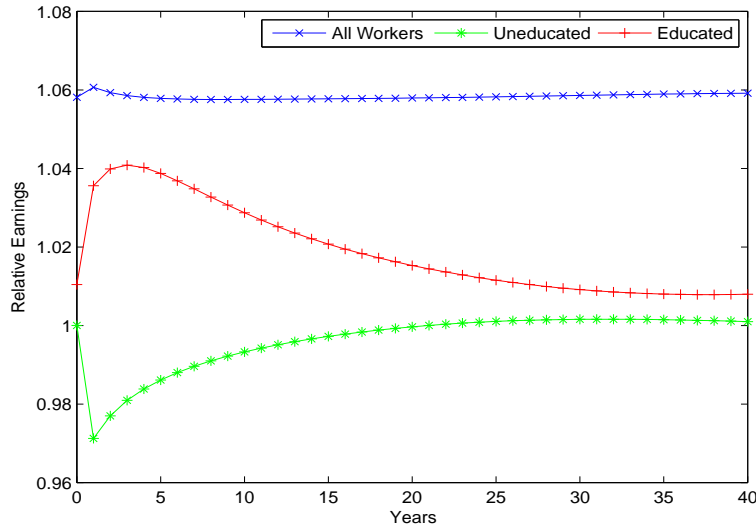
endogeneity of the skill supply in a Heckscher-Ohlin model with homogeneous agents. When agents are homogeneous, all workers must be indifferent between acquiring an education and remaining uneducated. Since there is a unique skill premium that ensures this indifference, the skill premium cannot change between the pre- and post-liberalization steady states, and as a consequence, the steady-state gains from trade are equal for all workers. In contrast, when workers are heterogeneous, as here, only the marginal worker is indifferent between pursuing an education and remaining uneducated. In order to increase the skill supply, the relative return to schooling must increase so as to induce less able individuals to pursue an education. As a result, even in the long run, educated workers gain from trade relatively more than uneducated workers despite the skill-supply adjustment.

Finally, Figure 4a also plots the evolution of per-period income inequality, defined as the per-period average income of an educated worker relative to that of an uneducated worker. Although income inequality is often used as a proxy for lifetime-earnings inequality, a more meaningful measure of actual inequality, the figure shows how per-period income inequality can misstate lifetime-earnings inequality. Workers make their decisions based on their lifetime earnings, and therefore high lifetime-earnings inequality tends to have an equalizing effect in the long run as it incentivizes uneducated workers to pursue an education. This augments the skill supply and decreases the returns to education. Per-period income inequality does not take this equalizing effect into account. It therefore overstates inequality when inequality is above its long-run level and understates inequality when it is below its long-run level. This highlights the importance of taking into account the endogeneity of the skill supply and its effect on future incomes when calculating inequality.

Inequality By Age

Figure 5 tracks the trajectory of the annualized lifetime earnings for the oldest workers relative to that of the youngest workers from the year before trade liberalization until forty years after its implementation. When all workers in these age groups are considered, inequality is higher than within education groups because the proportion of educated workers is higher among the oldest group than among the youngest. Because the proportion of educated workers is higher among the older age group, they, on average, gain more than younger workers immediately following trade liberalization. That there is little dynamics in the average within-age-group annualized lifetime-earnings inequality is a reflection of the fact that inequality within education groups moves in opposite directions for educated workers and uneducated workers.

Figure 5: Annualized Lifetime Earnings of Old Workers Relative to Young Workers



Notes: This figure plots relative annualized lifetime-earnings inequality between age cohorts within education groups from before trade liberalization until forty years into the transition.

In the steady states, there is no between-age-group inequality among uneducated workers because all uneducated workers receive the same wage. From the first period following trade liberalization their wages increase as the number of uneducated workers decreases. However, older workers do not live long enough to benefit from the entire recovery of low-skill wages and therefore they have lower annualized lifetime earnings than younger workers.

Educated workers experience the opposite effect. Older workers spend a greater proportion of

their remaining life working for the temporarily high wages than do younger workers. Younger workers remain in the workforce even after their wages have decreased in response to the trade-induced influx of educated workers. This effect causes an initial increase in the annualized lifetime-earnings inequality between old and young educated workers.

Unlike inequality among uneducated workers the inequality among educated workers does not reverse course immediately after the initial change. The reason is that the average ability of the workers induced by trade liberalization to acquire an education is lower than the average ability of workers who were already educated at the time. Since most newly educated workers are young, the increase in educated workers decreases the average ability of educated young workers relative to old workers. This has the effect of increasing the average annualized lifetime earnings of old educated workers relative to the average for young workers.

Table 4 reports the percentage changes, from the pre-liberalization steady state, in inequality between the youngest and oldest workers along the transition path. After the initial trade shock, average annualized lifetime earnings of old workers relative to young workers increases by 0.2%, while this can be decomposed into a 2.5% increase among educated workers and a 3.4% decrease for uneducated workers.

Table 4: Percentage Change in Annualized Lifetime Earnings for Oldest Relative to Youngest Workers

Group	Impact	5 years	10 years	25 years	Steady State
Uneducated	-3.4	-1.3	-0.6	0.1	0.0
Educated	2.5	2.8	1.8	0.1	0.0
Total	0.2	-0.1	0.0	0.0	0.1

Notes: Uneducated workers refers to workers who will never pursue an education.

In the long run, there is no change in inequality within education groups. However among all workers there is a 0.1% increase in the earnings of old workers relative to young workers. The reason is that there is a small increase in the college premium compared to before the trade liberalization. Since the proportion of educated workers is higher among older workers, they benefit more as a group from the higher college premium.

4.1.5 Importance of Skill Adjustment and Transition

How do the predictions for inequality and the gains from trade differ when the transition is not taken into account or when the endogeneity of the skill supply is ignored? The results reported in Table 5 demonstrate the importance of the endogenous skill supply and the transition.³⁸

Table 5: Percentage Changes Relative to Pre-Liberalization Steady State

	Impact	Steady State	Fixed Skill Supply
Uneducated Lifetime Earnings, 19-23	4.6	5.6	-0.7
Uneducated Lifetime Earnings, 64-68	1.5	5.6	-0.7
Educated Lifetime Earnings, 19-23	7.2	6.3	12.4
Educated Lifetime Earnings, 64-68	9.9	6.3	12.4
College Premium, All Ages	4.3	0.1	13.2

Notes: Uneducated workers refers to workers who will never pursue an education. The first four rows are the same as the gains from trade for these groups. The last row reports the relative lifetime earnings of an average educated worker relative to an uneducated worker.

Considering only the post-liberalization steady state leads to an understatement of the gains from trade for educated workers and an overstatement for uneducated workers. Therefore, ignoring the transition and focusing instead only on the long-run steady state leads to a substantial understatement of trade-induced inequality. Indeed, much of the inequality is realized during the transition as skill supplies adjust to the increased relative skill demand.

The situation is reversed when the endogeneity of the skill supply is not taken into account. Workers respond to the high returns to education by pursuing an education and thereby augmenting the skill supply, which reduces the returns to education. Thus, the endogenous skill supply tends to have an equalizing effect. Hence, if the skill adjustment is ignored, trade-induced inequality is overstated. In fact, were skill supplies fixed, uneducated workers would actually be worse off after trade liberalization, whereas when skill supplies are properly treated as endogenous they are better off, albeit by less than educated workers.

Finally, because all the dynamics in the model operate through changes in the skill supply, when the skill supply is fixed there are no dynamics. Therefore, ignoring either the transition or the


³⁸The equilibrium with fixed skill supply is computed by fixing the education cutoffs at their pre-liberalization values. Thus, the distribution of workers remains unchanged after trade liberalization along all three dimensions of the worker distribution, age, ability and educational attainment. This computation, as opposed to simply fixing the skill supply, allows me to maintain the age structure of the model.

endogeneity of the skill supply fails to account for the differential impact of trade across age groups. Old and young would be affected in the same way were there no transition.

4.1.6 Skill Supply Adjustment

In this section, I give a brief overview of the mechanisms that shape the evolution of the skill supply along the transition path. These mechanisms are important because the skill-supply adjustment is the source of all the dynamics in the model. Given the worker distribution in the steady state prior to trade liberalization, the trajectory of the labor supply is uniquely determined by the education cutoffs for uneducated workers. The education cutoffs, therefore, are the key to understanding the model's dynamics. Figure 6 shows the evolution of the education cutoffs for the youngest age group.

Figure 6: Education Cutoffs




Notes: The figure depicts the education cutoffs for the youngest workers. Older workers have similarly shaped paths but the levels are higher. The path is shown for eighty years into the transition to highlight the fluctuations in the education cutoff.

Workers who are old when the new trade policy is implemented will be undereducated relative to long-run levels as they made education decisions in their youth based on conditions prevailing prior to trade liberalization. This relative undereducation, which is reflected in Figure 6 by the relatively high pre-liberalization cutoff, causes a shortage in the relative supply of high-skill labor, which drives up the skill premium. The increase in the skill premium implies an increase in the

returns to education, which leads to a drop in the education cutoff. However, because of the scarcity of high-skill labor relative to its long-run level, the education cutoff drops to below its long-run level. Thus, the cohort of young workers at the time of trade liberalization will be overeducated relative to the long run. As older undereducated workers die and the younger overeducated workers take their place, the scarcity of high-skill labor gives way to a surplus. This process then reverses itself as the youth become undereducated relative to the long run in response to the surplus of high-skill labor. This cycle continues until the economy converges to the new post-liberalization steady state.

The economy experiences alternating periods of scarcity and surplus of high-skill labor relative to the long-run steady state. Figure 7b shows that the supply of high-skill labor overshoots its long-run level approximately twenty-five years following trade liberalization and then drops to below its long-run level twenty-five years after that. This cobweb effect is a result of the cohort structure of the model as well as workers' rational expectations and is not borne out of workers' irrational behavior or incorrect expectations about the general-equilibrium effect of changes in skill supply.³⁹ These cycles explain the oscillations, apparent in Figure 3, of the educated and uneducated workers' incomes.

Figure 7: Percentage Change in Skill Supply Along the Transition Path Relative to Pre-liberalization



Notes: The figure charts the transition of skill supply from one period before trade liberalization until sixty years into the transition.

The high-skill labor supply does not change between the pre-liberalization steady state and the first period following the liberalization. The reason is that the high-skill labor supply depends

³⁹This is in contrast to, for example, the cobweb model of Freeman (1976).

only on the distribution of educated workers, which cannot be changed within a period. Thus, the education decisions of workers only affects future, but not present, supply. High-skill labor supply only changes over time as a result of the education decisions of uneducated workers.

The low-skill labor supply, unlike the high-skill labor supply, is affected by both past and present education decisions. Past education decisions determine the current supply of uneducated workers. However, some of these workers will choose to become students and work only part time. Thus, for a given supply of uneducated workers, the supply of low-skill labor is decreasing in the proportion of uneducated workers who become students, or alternatively, increasing with the education cutoffs.

In the first period after trade liberalization, the number of uneducated workers is the same as in the steady state with the old trade policy as this number depends only on past decisions. However, because of the decrease in the education cutoffs, the supply of low-skill labor decreases in the first period after liberalization. After the first period, the number of uneducated workers decreases as more workers pursue an education. At the same time, the proportion of uneducated workers who become students decreases as the education cutoff gradually increases. These two opposing forces, the decrease in the number of uneducated workers and increase in the proportion of these workers who work full time, explain the hump in the low-skill labor supply that can be seen in the beginning of the transition in Figure 7a.

5 Conclusion

The objective of this paper has been to quantify the differential impact of trade liberalization on workers depending on their education level and age. To this end I develop a dynamic general-equilibrium trade model with endogenous education decisions. Two insights distinguish this paper from the existing trade literature. First, if trade changes relative skill demand and as a consequence also the skill premium, then workers will react to the change in the skill premium by adjusting their education or skill accumulation decisions. Second, the process of skill-supply adjustment to an increase in relative skill demand is far from instantaneous.

The model I study is rich enough to provide meaningful predictions on the impact of trade liberalization on workers along several dimensions. At the same time, the model remains tractable

enough to allow for a solution along the economy's entire transition path following trade liberalization. Thus, the model sheds light not only on the effect of trade liberalization on workers by education and age, but also by birth cohort.

After calibrating the model to 2007 US data, I simulate the economy's equilibrium transition path in response to a once-and-for-all removal of policy trade barriers. I find that the gains from trade are most unevenly divided among workers alive at the time of trade liberalization. This is true both if we consider workers by education within age groups or vice versa. In particular, for this generation of workers, educated workers gain more the older they are, with the oldest educated workers gaining 9.9% and the youngest gaining 7.2%. Among uneducated workers, the reverse is true, and as a result, the oldest uneducated workers gain the least from trade liberalization, a mere 1.5%. Among workers in future generations, born subsequent to trade liberalization, the gains from trade are divided much more equitably.

I find that trade leads to a substantial increase in inequality, although most of this inequality is transitory. The college premium, defined as the discounted lifetime earnings of educated workers relative to that of uneducated workers, increases initially by 4.3% when all age groups are considered. However, this can be decomposed into an 8.2% increase for the oldest workers and a 2.6% increase for the youngest workers. In addition, I find that, even within education groups, annualized lifetime earnings differ by age. All these measures of inequality are highest immediately following trade liberalization and decrease during the economy's transition. Nevertheless, even in the long run the college premium is higher than pre-liberalization, but by only 0.2 percentage points.

The reason for the decrease in inequality during the transition is that the high returns to education, the cause of the inequality, induce workers to pursue an education. This leads to an increase in the supply of skilled labor which, over time, lowers the returns to education.

The upshot of these results is that the dynamics due to the endogeneity of the skill supply are essential for understanding the full impact of trade on workers. On the one hand, ignoring the transition leads to an understatement of trade-induced inequality as most of the inequality is transitory. On the other hand, not allowing skill supplies to adjust to trade liberalization leads to an overstatement of trade-induced inequality as it does not account for the equalizing effect of the endogenous increase in the skill supply.

One of the main lessons of my results is that accounting for changes in the educational composition of the workforce is essential for understanding the full implications of trade liberalization. In particular, this paper highlights that education plays a crucial role in mitigating the adverse distributional impact of trade liberalization, at least in the long run. This has potentially important implications for policy makers tasked with devising policies aimed at combating the unequalizing effects of trade liberalization. Indeed, policy makers have long recognized the important role education can play in complementing trade-liberalizing policies. One example of this recognition is the Trade Adjustment Assistance program in the United States created under the Trade Act of 1974, which provides training assistance to workers negatively impacted by trade. The results in this paper can help inform precisely these types of trade-assistance policies.

These programs also highlight a path for future research. Many training-assistance programs are predicated on the assumption that some workers will not pursue an education or further their training in response to changes in trade policy, even if it is desirable for them to do so. Credit constraints and financial frictions are likely culprits, as these make it difficult for workers to invest in an education. Understanding how such frictions affect inequality and skill accumulation decisions can help policy makers design better training-assistance programs. As such, incorporating these frictions into the model presented in this paper is a promising avenue for future research.

References

- Anderson, J. E. and E. van Wincoop (2003). Gravity with gravitas: A solution to the border puzzle. *The American Economic Review* 93(1), 170–192.
- Anderson, J. E. and E. van Wincoop (2004). Trade costs. *Journal of Economic Literature* 42, 691–751.
- Artuç, E. (2009). Intergenerational effects of trade liberalization.
- Artuç, E., S. Chaudhuri, and J. McLaren (2010). Trade shocks and labor adjustment: A structural empirical approach. *The American Economic Review* 100(3), 1008–1045.
- Atkeson, A. and A. Burstein (2007). Innovation, firm dynamics, and international trade. Technical report, National Bureau of Economic Research.
- Bernard, A. B., J. Eaton, J. B. Jenson, and S. Kortum (2000). Plants and productivity in international trade. Technical report, National Bureau of Economic Research.
- Bernard, A. B. and J. B. Jensen (1997). Exporters, skill upgrading, and the wage gap. *Journal of International Economics* 42(1), 3–31.
- Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2007). Firms in international trade. *The Journal of Economic Perspectives* 21(3), 105–130.
- Burstein, A. and J. Vogel (2010). Globalization, technology, and the skill premium: A quantitative analysis. Technical report, National Bureau of Economic Research.
- Bustos, P. (2007). The impact of trade on technology and skill upgrading evidence from argentina. *Working Papers (Universitat Pompeu Fabra. Departamento de Economía y Empresa)* (1189), 1.
- Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *The American Economic Review* 101(1), 304–340.
- Campos-Vázquez, R. M. (2010). Why did wage inequality decrease in mexico after nafta? Technical report, El Colegio de México, Centro de Estudios Económicos.
- Chaney, T. (2008). Distorted gravity: The intensive and extensive margins of international trade. *The American Economic Review* 98(4), 1707–1721.

- Cosar, K. (2010). Adjusting to trade liberalization: Reallocation and labor market policies. *V manuscript*.
- Costantini, J. and M. Melitz (2007). The dynamics of firm-level adjustment to trade liberalization. *The Organization of Firms in a Global Economy*, 107–141.
- Davis, D. R. and J. Harrigan (2007). Good jobs, bad jobs, and trade liberalization.
- Egger, H. and U. Kreickemeier (2009). Firm heterogeneity and the labor market effects of trade liberalization. *International Economic Review* 50(1), 187–216.
- Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011). Income distribution, product quality, and international trade. *Journal of Political Economy* 119(4), 721–765.
- Falvey, R., D. Greenaway, and J. Silva (2010). Trade liberalisation and human capital adjustment. *Journal of International Economics* 81(2), 230–239.
- Findlay, R. and H. Kierzkowski (1983). International trade and human capital: A simple general equilibrium model. *The Journal of Political Economy*, 957–978.
- Freeman, R. B. (1976). A cobweb model of the supply and starting salary of new engineers. *Industrial and Labor Relations Review* 29(2), 236–248.
- Goldberg, P. K. and N. Pavcnik (2007). Distributional effects of globalization in developing countries. *Journal of Economic Literature* 45, 39–82.
- Harrigan, J. and A. Reshef (2011). Skill biased heterogeneous firms, trade liberalization, and the skill premium. Technical report, National Bureau of Economic Research.
- Harris, R. G. and P. E. Robertson (2013). Trade, wages and skill accumulation in the emerging giants. *Journal of International Economics* 89(2), 407–421.
- Harrison, A., J. McLaren, and M. McMillan (2012). Recent perspectives on trade and inequality. *Annual Review of Economics* 3(1), 261–289.
- Heckman, J. J., L. Lochner, and C. Taber (1998). Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents. *Review of Economic Dynamics* 1(1), 1–58.

- Helpman, E., O. Itzhoki, and S. Redding (2010). Inequality and unemployment in a global economy. *Econometrica* 78(4), 1239–1283.
- Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *The Quarterly Journal of Economics* 107(1), 35–78.
- Kim, D.-I. and R. H. Topel (1995). Labor markets and economic growth: Lessons from korea’s industrialization, 1970-1990. In *Differences and Changes in Wage Structures*, pp. 227–264. University of Chicago Press.
- Kwark, N. and C. Rhee (1993). Educational wage differentials in korea.
- Lileeva, A. and D. Trefler (2010). Improved access to foreign markets raises plant-level productivity. . . for some plants. *The Quarterly Journal of Economics* 125(3), 1051–1099.
- Matsuyama, K. (2007). Beyond icebergs: Towards a theory of biased globalization. *The Review of Economic Studies* 74(1), 237–253.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695–1725.
- Melitz, M. J. and S. J. Redding (2013). Firm heterogeneity and aggregate welfare. Technical report, National Bureau of Economic Research.
- Porto, G. G. (2006). Using survey data to assess the distributional effects of trade policy. *Journal of International Economics* 70(1), 140–160.
- Reed, W. J. (2001). The pareto, zipf and other power laws. *Economics Letters* 74(1), 15–19.
- Robertson, R. (2007). Trade and wages: Two puzzles from mexico. *The World Economy* 30(9), 1378–1398.
- Topel, R. H. (1997). Factor proportions and relative wages: The supply-side determinants of wage inequality. *The Journal of Economic Perspectives* 11(2), 55–74.
- Tybout, J. R. Plant-and firm-level evidence on “new” trade theories. *Integration & Technology*, 67.
- Verhoogen, E. A. (2008). Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. *The Quarterly Journal of Economics* 123(2), 489–530.

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

A Theory Appendix

Technology and Export Cutoffs

In this section, I describe the algorithm I use for computing the technology and export cutoffs. The technology used by the least productive firms is $v = 0$ which ensures zero profits. Given this starting point, the remaining cutoffs are computed using the following procedure starting with $v = 1$ and $\tilde{v} = 0$:

1. Find the productivity φ_v at which a firm is indifferent between using technology v and technology \tilde{v} , which is the highest technology used lower than v .
2. Calculate the profit at φ_v for every technology $v' > v$.
 - (a) If there exists a v' such that $\pi(\varphi_v, v') > \pi(\varphi_v, v)$, then technology v is not used and return to step 1 with $v + 1$.
 - (b) Otherwise, technology v is used and the adoption cutoff is φ_v . If φ_x has already been found then return to step 1 with $v + 1$.
3. Denote φ_{vx} the productivity at which a firm would be indifferent between exporting and selling only domestically if it uses technology v .
 - (a) If $\varphi_{vx} < \varphi_v$, then $\varphi_x = \varphi_{vx}$.
 - (b) Otherwise, if $\varphi_{vx} \leq \varphi_v$, then $\varphi_x = \varphi_v$.
 - (c) Otherwise, φ_x has still not been reached, and return to step 1 with $v + 1$.

Worker Policy

Consider first a worker in period t described by the triplet $(a, b, 1)$, that is, an educated worker with ability a in age group b . Such a worker's value is her current wage plus the discounted expected continuation value of being educated in period $t + 1$,

$$V_t(a, b, 1) = aw_{ht} + \beta \tilde{V}_{t+1}(a, b, 1),$$

where \tilde{V} is the expected value of the worker in the subsequent period. The expectation takes into account that the worker may age or die, and therefore,

$$E \left[\tilde{V}_t(a, b, j) \right] = (1 - \delta_b - D_b) V_t(a, b, j) + \delta_b V_t(a, b + 1, j),$$

for $j \in \{0, 1\}$ and where $V_t(a, B+1, j) = 0$, since $B+1$ corresponds to death.

The value of working full time for an uneducated worker in period t with ability a in age group b is

$$V_t^w(a, b) = w_{\ell t} + \beta \tilde{V}_{t+1}(a, b, 0).$$

This worker's value of pursuing an education is

$$V_t^s(a, b) = -(\ell_e w_{\ell t} + h_e w_{ht}) + m w_{\ell t} + \beta \left[(1 - \delta_e) \tilde{V}_{t+1}(a, b, 0) + \delta_e \tilde{V}_{t+1}(a, b, 1) \right],$$

where the first two terms are the worker's income net of tuition costs and the remaining terms are the worker's continuation values accounting for the fact that the student will become educated with only a δ_e probability. The value of an uneducated worker with ability a in age group b is the maximum of the value of working full time and becoming a student

$$V_t(a, b, 0) = \max \{V_t^w(a, b), V_t^s(a, b)\}.$$

From the above equations it is apparent that if worker values are known for a given future period T , then the above formulas allow for the backward induction of worker values for all periods $t \leq T$ for a given path of wages. In practice, I will calculate worker values by backward inducting from the terminal steady state of the economy. I will therefore show how worker values are computed in a steady state.

Suppose that the steady state wages are w_ℓ and w_h . The expected discounted remaining lifetime of a worker in age group b is $\lambda_b = (1 + \beta \delta_b \lambda_{b+1}) / (1 - \beta(1 - \delta_b - D_b))$ with the initial condition $\lambda_{B+1} = 0$. Therefore, in steady state, $V(a, b, 1) = \lambda_b w_h$. If an uneducated worker decides to work full time in age group b , then she will also choose to do so when she reaches older age groups because the education cutoffs, $\{a_b\}_{b=1}^B$, are increasing with age. For age group B , if $a < a_B$, then $V(a, B, 0) = \lambda_B w_\ell$. If instead $a > a_B$, then

$$V(a, B, 0) = \frac{-(\ell_e w_\ell + h_e w_h) + m w_\ell + \delta_e \lambda_B w_h}{1 - \beta \delta_e (1 - \delta_B - D_B)}.$$

Given the values for age group B , backwards induction yields

$$V(a, b, 0) = \begin{cases} \lambda_b w_\ell & a \leq a_b \\ \frac{-(\ell_e w_\ell + h_e w_h) + m w_\ell + \beta [\delta_e \tilde{V}(a, b, 1) + (1 - \delta_e) \delta_b V(a, b+1, 0)]}{1 - \beta \delta_e (1 - \delta_B - D_B)} & a > a_b \end{cases}.$$

This completes the calculation of worker values for the steady state.

Worker Distribution

Suppose the worker distribution in period t is $W_t(a, b, e)$ and the education cutoffs are $\{a_b\}_{b=1}^B$. It will be useful to define an auxiliary distribution, $\tilde{W}_t(a, b, \tilde{e})$, where $\tilde{e} = 0$ if the worker is uneducated and works full time, $\tilde{e} = 1$ if the worker is educated and $\tilde{e} = 2$ if the worker is pursuing an education in the given period. This distribution describes the worker distribution given their optimization decisions in period t . For educated workers these two distributions are identical, $W_t(a, b, 1) = \tilde{W}_t(a, b, 1)$. For uneducated workers the education cutoffs govern the relationship between W_t and \tilde{W}_t

$$\tilde{W}_t(a, b, 0) = \begin{cases} W_t(a, b, 0) & a \leq a_{bt} \\ 0 & a > a_{bt} \end{cases}, \quad \tilde{W}_t(a, b, 2) = \begin{cases} W_t(a, b, 0) & a > a_{bt} \\ 0 & a \leq a_{bt} \end{cases}.$$

The measure of uneducated workers for any (a, b) in period $t+1$ is the sum of two groups, workers who pursued an education in period t , but failed to become educated and uneducated workers who worked full-time in period t . Each of these groups, in turn, is a sum of workers who remained in age group b and those who aged from age group $b-1$,

$$W_{t+1}(a, b, 0) = \bar{W}_{t+1}(a, b, 0) + (1 - \delta_e) \bar{W}_{t+1}(a, b, 2),$$

where $\bar{W}_{t+1}(a, b, j) = (1 - \delta_b - D_b) \tilde{W}_t(a, b, j) + \delta_b \tilde{W}_t(a, b-1, j)$ for $j \in \{0, 1, 2\}$. Denoting the measure of newborns in each period by N , $\delta_b \tilde{W}_t(a, 0, 0) = Ng_w(a)$ is the measure of newborns with ability a in each period. Since newborns are born uneducated and did not pursue an education in the previous period, $\tilde{W}_t(a, 0, j) = 0$ for $j \in \{1, 2\}$.

The distribution of educated workers in period $t+1$ is calculated analogously to the distribution of uneducated workers. The difference being that educated workers in $t+1$ are the sum of workers who were already educated in period t and students in period t who succeeded in becoming educated,

$$W_{t+1}(a, b, 1) = \bar{W}_{t+1}(a, b, 1) + \delta_e \bar{W}_{t+1}(a, b, 2).$$

Finally, the measure of newborns, N , in each period is constant and ensures that the measure of births exactly equals the measure of deaths in each period. This condition, with population normalized to one, means that the measure of newborns is given by the recursive formula

$$N = \left(\sum_{b=1}^B \zeta_b \right)^{-1}, \quad \zeta_b = \zeta_{b-1} \left(\frac{\delta_{b-1}}{\delta_b + D_b} \right) \quad \text{and} \quad \zeta_1 = \frac{1}{\delta_1}.$$

B Calibration Appendix

Bins

The CPS divides the firm size space into 6 bins. The smallest bin contains firms with between one and nine workers. The remaining bins are 10-24, 25-99, 100-499, 500-999 and 1,000 or more workers.

The Economic Census, beginning in 2007, divides the size space into twenty-five bins. Up to firms with fifty workers, the bins are divided by increments of five workers. The remaining bins have the following lower boundaries: 75, 100, 150, 200, 300, 400, 500, 750, 1,000, 1,500, 2,000, 2,500, 5,000 and 10,000.

Because firms with 10,000 or more workers represent approximately 40% of total employment in the population, it is reasonable to assume that also those firms are able to switch technologies. I therefore supplement the Economic Census data with COMPUSTAT data in order to divide the bin with 10,000 or more workers into 8 more bins with lower boundaries: 10,000, 25,000, 50,000, 75,000, 100,000, 250,000, 500,000 and 1,000,000. I choose 1,000,000 as the upper boundary because that is approximately the size of the largest firm in the United States.

The COMPUSTAT database includes only publicly traded firms and the size information includes worldwide employment. I adjust for the under-representation of firms in COMPUSTAT by weighting the COMPUSTAT firms by a factor that ensures that the number of total firms with more than 10,000 workers is equal to the number reported by the Economic Census. I make a similar adjustment for firm size, so that aggregate employment in the 10,000 and more workers bin is equal to the total employment in this bin as reported in the Economic Census data.

Size Adjustment

The firm size data reports numbers of workers. However, the model contains only predictions about the number of labor units hired by a firm. This is because firms can, for example, hire two

units of high-skill labor either by hiring one educated worker with innate ability equal to two or two educated workers with innate ability equal to one. The number of low-skill labor units is the total wage paid to uneducated workers divided by the average wage of uneducated workers, and an analogous computation holds for high-skill labor. I adjust the CPS bin boundaries by a factor equal to total labor units in the CPS bin divided by the total number of workers in the CPS bin. This factor is not equal across CPS bins as larger firms tend to pay higher wages to workers with similar educational attainment.

B.1 Calibration of Firm-Level Technologies

Let lb_v be the lower boundary of size bin v , and let s_v be the smallest firm in size bin v .⁴⁰ It must be the case that $lb_v \leq s_v$, and if the inequality is strict, then there must be a discrete jump in size from the largest firm in bin $v - 1$ to the smallest firm in bin v . The idiosyncratic productivity of a firm with size s_v is, on account of the Pareto distribution of φ ,

$$\varphi_v = [1 - \Pr(\text{size} < lb_v)]^{-\frac{1}{\theta_f}}.$$

Because I have data on the number of firms in each size bin, the proportion of firms smaller than s_v or lb_v , $\Pr(\text{size} < lb_v)$, is data. Therefore, the model together with the data imply a unique idiosyncratic productivity, φ_v , for the marginal technology upgrader for every v .

For each CPS bin I have data on relative skill expenditure. Suppose the CPS bins were partitioned in the same way as the size bins and let x_v be the relative low- to high-skill labor expenditure in bin v . In that case, setting the relative high-skill labor requirement for technology v to $\lambda_v = x_v w_h / w_\ell$ would ensure that the relative expenditure in each bin would match the data since, by assumption, all firms in the bin employ the same technology. However, the size bins are, in fact, a nontrivial partition of the CPS bins, and I do not, therefore, have direct data on λ_v .

In order to identify the set $\{\lambda_v\}_{v=1}^{32}$, I start by interpolating, with a cubic spline, the values for λ at the CPS bin boundaries.⁴¹ I assume that within CPS bin j the relationship between s_v , the size of the smallest firm that adopts technology v , and λ_v , the relative high-skill labor requirement

⁴⁰The sizes of firms in each bin are adjusted to be in terms of labor units rather than workers to be consistent with the model, as described above.

⁴¹I use the average firm size in each CPS bin and the average skill intensity within the CPS bin as the datapoints for the interpolation.

of technology v , is quadratic,

$$\lambda_v(s_v) = a_j(s_v - lb_j)^2 + b_j(s_v - lb_j) + c_j,$$

where λ_j is the interpolated value of λ at the lower boundary of the CPS bin j and lb_j is the lower boundary of this CPS bin. Because the bin endpoints are determined by the interpolation, this quadratic equation has one degree of freedom. I calibrate the technologies within each CPS bin given coefficients (a_j, b_j, c_j) , so that s_v implies a λ_v , and use the degree of freedom to match the data on relative skill demand in each CPS bin.⁴²

Identification of Technology Productivity and Fixed Cost

I calibrate the productivity of the technologies, $\{\alpha_v\}_{v=1}^{34}$, and fixed costs of technology usage, $\{f_v\}_{v=1}^{34}$, directly from the firms' optimization conditions. Indeed, this is one of the strengths of my calibration methodology. I begin by calibrating technology $v = 1$ used in the first size bin and progressively move along the size bins until all technologies are calibrated.

Consider the calibration of any technology v , other than the technology adopted by the marginal exporter, which I will consider shortly. I assume that the size of the smallest firm in bin v is equal to the lower boundary of this size bin, $lb_v = s_v$. This implies a λ_v for coefficients (a_j, b_j, c_j) . In addition, the size of smallest firm in bin v must equal the labor usage per unit produced, $(1 + \lambda_v) / (\varphi_v \alpha_v)$ times the quantity produced by the firm. This equality together with the optimal quantity produced by the firm derived in Section 2 yields,

$$\alpha_v = \frac{1}{\varphi_v} \left[\frac{R\rho^\sigma (1 + \mathbf{1}_{q_F > 0} \tau^{1-\sigma}) (w_\ell + \lambda_v w_h)^{-\sigma} (1 + \lambda_v)}{s_v} \right]^{\frac{1}{1-\sigma}}.$$

All the values on the RHS of this equation have known values. This value of α_v is the calibrated value of the productivity of technology v .

Since s_v is the smallest firm that takes on technology v , firm optimization implies that an s_v sized firm earns the same profit if it employs technology $v - 1$ as if it employs technology v . Because

⁴²The lower boundary of the CPS bin gives $\lambda_v(lb_v) = \lambda_j$ which implies that $c = \lambda_j$. The upper boundary of the CPS bin gives $\lambda_{j+1} = a_j(lb_{j+1} - lb_j)^2 + b_j(lb_{j+1} - lb_j) + \lambda_j$ which implies that $a_j = [b_j(lb_{j+1} - lb_j) + \lambda_j - \lambda_{j+1}] / (lb_{j+1} - lb_j)^2$. Thus, I match the data on relative skill demand within in CPS bin j by finding the appropriate a_j .

technology v is not the one adopted by the marginal exporter, adoption of this technology does not induce the firm to change its export status. Therefore,

$$f_v = \frac{(A_H + \mathbf{1}_{q_F > 0} A_F)}{\varphi^{1-\sigma}} \left[\left(\frac{w_\ell + \lambda_v w_h}{\alpha_v} \right)^{1-\sigma} - \left(\frac{w_\ell + \lambda_{v-1} w_h}{\alpha_{v-1}} \right)^{1-\sigma} \right] + f_{v-1}.$$

Since the technologies are calibrated in order, α_{v-1} and f_{v-1} have already been calibrated at this point. Therefore, all the parameters on the RHS are known. Thus, this equation provides the calibrated value of f_v . If $v = 1$, then a zero profit condition is used. In other words, the least productive active firm must earn zero profits which is equivalent to imposing $f_0 = \alpha_0 = 0$ in the above equation.

Marginal Exporter and Identification of Fixed Export Cost

Consider now the calibration of the technology adopted by the marginal exporter, denoted by v_x . I assume that the marginal exporter also upgrades technology. Bernard, Jensen, Redding and Schott (2007) report that approximately 4% of US firms export. This value roughly corresponds to the proportion of firms with more than fifty workers.⁴³ I therefore assume that the marginal exporter is the smallest firm in the size bin consisting of firms employing between fifty and seventy-five workers which also means that $v_x = 11$.

Technology v_x cannot be calibrated with the same methodology as the other technologies. The reason is that the assumption that s_{v_x} is equal to the lower bin boundary, i.e. fifty, as was assumed for the other size bins, creates an inconsistency between the model and the data. The model implies that there is a discrete jump in firm size at the marginal exporter. If the smallest exporter employs fifty workers it would imply that the next smallest firm, i.e. the largest non-exporter, employs fewer than forty-five workers. However, this would imply that there were no firms in the 45-50 worker size bin which is at odds with the data. In this case, therefore, $s_{11} > lb_{11} = 50$ and I treat s_{11} as an unknown.

As with the other technologies firm optimization implies that

$$s_{11} = \frac{R\rho^\sigma (1 + \tau^{1-\sigma}) (w_\ell + \lambda_{11} w_h)^{-\sigma} (1 + \lambda_{11})}{(\alpha_{11} \varphi_x)^{1-\sigma}}. \quad (1)$$

However, unlike in the calibration of the other technologies, this equation has two unknowns, s_{11}

⁴³ Approximately 3.9% of firms employ 50 or more workers.

and α_{11} . I assume that the increase in the cost efficiency from $v_x - 1$ to v_x , is 4% which roughly corresponds to the percent which starting exporters are more productive than non-exporters as reported in Bernard and Jensen (2004).⁴⁴ This implies that

$$\frac{\alpha_{11}}{w_\ell + \lambda_{11}w_h} = 1.04 \frac{\alpha_{10}}{w_\ell + \lambda_{10}w_h}. \quad (2)$$

The solution to equations 1 and 2 yield values for α_{11} and s_{11} .

The remaining parameters to calibrate are f_x and f_{v_x} , the fixed export cost and the fixed cost of using technology v_x . The marginal exporter must earn the same profit when exporting and upgrading to technology v_x as when it serves only the domestic market and uses technology $v_x - 1$. This optimization condition implies that

$$f_x + f_{v_x} = \frac{(A_H + A_F)}{\varphi_x} \left(\frac{w_\ell + \lambda_{v_x}w_h}{\alpha_{v_x}} \right)^{1-\sigma} - \frac{A_H}{\varphi_x} \left(\frac{w_\ell + \lambda_{v_x-1}w_h}{\alpha_{v_x-1}} \right)^{1-\sigma} + f_{v_x-1}.$$

Because the marginal exporter upgrades technology and switches export status the indifference condition only allows the identification of the sum of the fixed costs.

The complementarity between exporting and technology upgrading, and the fact that the marginal exporter makes a joint decision to both upgrade technology and switch export status implies that it is unprofitable for the firm to do one without the other. This condition imposes lower bounds on f_x and f_{v_x} . Specifically, the firm must prefer to serve only the Home market if it does not upgrade technology, and therefore

$$f_x > \frac{A_F}{\varphi_x} \left(\frac{w_\ell + \lambda_{v_x-1}w_h}{\alpha_{v_x-1}} \right)^{1-\sigma} + f_{v_x-1} = \underline{f_x}.$$

Similarly, the firm must prefer to use technology $v_x - 1$ over v_x if it does not export. This implies

$$f_{v_x} > \frac{A_H}{\varphi_x} \left[\left(\frac{w_\ell + \lambda_{v_x}w_h}{\alpha_{v_x}} \right)^{1-\sigma} - \left(\frac{w_\ell + \lambda_{v_x-1}w_h}{\alpha_{v_x-1}} \right)^{1-\sigma} \right] + f_{v_x-1} = \underline{f_{v_x}}.$$

The calibration implies that, $f_x \in (\underline{f_x}, \underline{f_x} + x)$ and $f_{v_x} \in (\underline{f_{v_x}}, \underline{f_{v_x}} + x)$, where $x = f_x + f_{v_x} - \underline{f_x} - \underline{f_{v_x}}$ represents the leeway the calibration provides in identifying each of the two fixed costs.

The result of the calibration is $x = 10,510$, which implies that the fixed exporting cost is in the interval (98,425, 108,935). These numbers are in 2007 US dollars.

⁴⁴This is not exactly the same as the increase in the productivity of the marginal exporter in my model. Simulations with alternative numbers indicate that the results are not sensitive to different choices of the increase in productivity for the marginal exporter.

B.2 Calibration Results

The parameters, other than the firm-level technological parameters, used in the baseline simulation are summarized in Table 6.

Table 6: Calibrated Parameters

Parameter	Value	Moment
β	0.97	average annual real interest rate of 3%
B	10	workers are divided into 10 age groups
δ_b	0.2	age group corresponds to a 5-year window
δ_e	0.22	average of 4.5 years to complete BA
σ	3.8	estimate in Bernard et al. (2003)
τ	1.91	exporters ship 14% of shipments abroad
θ_f	9.3	relative skill supplies
θ_w	2.92	ratio of educated workers to high-skill labor
f_x	\$103,680	calibration of technology menu
$\{D_1, \dots, D_{10}\}$	$0.01 * \{0, 0.1, 0, 0, 0.2, 0.3, 0.6, 1.6, 4.6, 12.4\}$	decrease in labor market participation by age

C Simulation Appendix

In this Appendix, I provide a brief description of the algorithm I use for computing an equilibrium path from any initial distribution of workers to a terminal steady state. In practice, I solve for an equilibrium path between an initial steady state and a terminal steady state, but the algorithm requires knowledge only of the worker distribution in the initial steady state and the worker values in the terminal steady state.

To solve for a steady state, I solve for a triplet (w_ℓ, w_h, R) . This triplet must satisfy the equilibrium conditions laid out in the main text. Solving the initial and terminal steady states gives the initial worker distribution and terminal worker values, as described in Appendix A, required to compute the equilibrium transition path. The algorithm I will describe is for a given transition length T . The transition length is chosen to be long enough to ensure convergence to the terminal steady state.

1. Guess a path of wages $\{w_{\ell t}, w_{ht}\}_{t=1}^T$.
2. Compute optimal worker policy given $\{w_{\ell t}, w_{ht}\}_{t=1}^T$ by solving backward for worker values and education cutoffs as described in Appendix A.
3. Using worker policy computed in 2 and the initial worker distribution, solve forward for the path of skill supplies and measure of students, $\{L_t, H_t, S_t\}_{t=1}^T$, as described in Appendix A and the main text.
4. For each t , given (L_t, H_t, S_t) , solve for an equilibrium, i.e. find $(\tilde{w}_{\ell t}, \tilde{w}_{ht}, R_t)$ such that firms optimize and all markets clear.
5. If the algorithm has not yet converged then adjust the guess for the path of wages and return to Step 2 and iterate until convergence. The convergence criteria is $\max\{(w_{\ell t} - \tilde{w}_{\ell t})/w_{\ell t}, (w_{ht} - \tilde{w}_{ht})/w_{ht}\} < 10^{-7}$.

I adjust the guesses in each iteration so that for each period the new guess is a linear combination of w_{it} and \tilde{w}_{it} for $i \in \{\ell, h\}$ with a weight of 0.05 on \tilde{w}_{it} . A higher weight speeds up the convergence, but makes the convergence unstable.