From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2014

Essays in Trade and Labor Economics

Danziger, Eliav

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01x346d640h

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Essays in Trade and Labor Economics

Eliav Danziger

A Dissertation

Presented to the Faculty

of Princeton University

in Candidacy for the Degree

of Doctor of Philosophy

Recommended for Acceptance

by the Department of

Economics

Adviser: Professor Esteban Rossi-Hansberg

September 2014

© Copyright by Eliav Danziger, 2014.

All rights reserved.

Abstract

This collection consists of three chapters whose overarching theme is inequality. The first chapter explores the differential impact of trade liberalization on workers according to their educational attainment when schooling decisions are endogenous. My main findings are that the oldest educated workers gain most from trade liberalization, while the oldest uneducated workers gain the least. Moreover, I find that ignoring the economy's transition path following trade liberalization leads to an understatement of trade-induced inequality, while ignoring the endogeneity of workers' schooling decisions leads to an overstatement of trade-induced inequality. The second chapter explores the differential impact of global changes on workers depending on their sectoral choice. I develop a model with endogenous unemployment, labor-force-participation and sectoral-entry decisions. In order to quantify the impact of global changes on the Canadian economy, I calibrate the model to Canadian data from 1993-2010. My main findings are that world price changes lead to an increase in unemployment in the manufacturing and service sectors, while they lead to a decrease in unemployment in the the natural resource and agriculture sectors. Labor-force participation increases in the natural resource sector as a result of changing world prices, but changes little in other sectors. Finally, I find that changes in world prices from 1993 to 2010 lead to a 7.5% decrease in aggregate welfare in the short run and to a 5.5% decrease in the long run. The final chapter, cowritten with Leif Danziger, shows that a graduated minimum wage, in contrast to a constant minimum wage, can provide a strict Pareto improvement over what can be achieved with an optimal income tax. The reason is that a graduated minimum wage requires high-productivity workers to work more to earn the same income as low-productivity workers, which makes it more difficult for the former to mimic the latter. In effect, a graduated minimum wage allows the low-productivity workers to benefit from second-degree price discrimination which increases their income.

Acknowledgements

I am indebted to my advisor, Esteban Rossi-Hansberg, for his encouragement and support. His helpful advice and suggestions have been invaluable. Although formally they were not my advisors, Gene Grossman's and Stephen Redding's doors were always open to me. Their vision of what it means to be a scholar has deeply influenced my own. Without their guidance, my research, past, present and future, would undoubtedly have suffered.

Oleg Itskhoki, Eduardo Morales and Richard Rogerson all provided excellent suggestions and invaluable discussion. I would also like to thank the International Economics Section of Princeton University both for the funding they provided and for creating a supportive forum for discussing my research and learning from others' research. In particular, I would like to thank the many participants in the IES student workshop, both students and faculty, for creating an environment in which my research could thrive.

Sharon Ernst and Laura Hedden provided excellent administrative support. Without their dedication, neither the graduate program nor the IES would function.

Although I am deeply indebted to all of those already mentioned, their contribution pales in comparison to that of my parents, Leif and Ruhama. The only thing that they did not provide for me is the skill to adequately express my gratitude for all that they have done. But how could they? The language has yet to be invented that can capture even a fraction of my appreciation and gratitude.

Contents

Abstract iii

Acknowledgements iv

List of Tables vi

List of Figures vii

1 Skill Acquisition and the Dynamics of Trade Induced Inequality 1

1.1 Introduction 2

1.2 The Model 10

1.2.1 Setup 10

1.2.2 Firms 11

1.2.3 Workers 18

1.2.4 Equilibrium 22

1.3 Calibration 23

1.3.1 Data 23

1.3.2 Calibration Details 25

1.3.3 Discussion of Calibration 31

1.4 Quantitative Analysis 33

1.4.1 Results 34

1.5 Conclusion 50

2 The Impact of World Price Changes on a Small Open Economy 53

2.1 Introduction 54
2.2 Model 59
2.2.1 Setup 59
2.2.2 Workers 60
2.2.3 Equilibrium 66
2.3 Mechanisms 67
2.3.1 Unemployment and Labor-Market Participation 68
2.3.2 Compositional Effects 70
2.3.3 General Equilibrium Effects 71
2.4 Methodology 73
2.4.1 Identification of Prices 75
2.5 Calibration 77
2.5.1 Human Capital 78
2.5.2 Sectoral Productivities 80
2.5.3 Job-Opportunity Distribution 82
2.5.4 Sector-Entry Cost Distribution 83
2.5.5 Calibration Results and Discussion 84
2.5.6 Initial Worker Distribution 86
2.6 Quantitative Analysis 87
2.6.1 Prices 89
2.6.2 Sectoral Choice and General Equilibrium Effects 90
2.6.3 Compositional Effects 92
2.6.4 Labor-Market Outcomes 94
2.7 Welfare 100
2.8 Conclusion 103

3 A Pareto-Imroving Minimum Wage **106**

3.1 Introduction 106

3.2 The Standard Competitive Model with Taxes 109
3.3 A Graduated Minimum Wage 111
3.4 A Continuous Piecewise Differentiable Graduated Minimum Wage 120
3.5 Heterogeneous Firms 122
3.6 Conclusion 126

A Appendix to Chapter 1 **128**
A.1 Theory Appendix 128
A.2 Calibration Appendix 132
A.2.1 Calibration of Firm-Level Technologies 133
A.2.2 Calibration Results 138
A.3 Simulation Appendix 139

B Appendix to Chapter 2 **140**
B.1 Calibration Appendix 140
B.1.1 Worker Distribution 140
B.2 Numerical Implementation 142
B.2.1 Worker Values and Policy 142
B.2.2 Worker Distribution 144
B.2.3 Simulation 145

C Appendix to Chapter 3 **148**

List of Tables

1.1 Gains from Trade for Workers Alive at Implementation of Trade Liberalization 38

1.2 Gains for Workers Born After Trade Liberalization by Birth Cohort 39

1.3 Percentage Change in the College Premium After Trade Liberalization . . . 42

1.4 Percentage Change in Annualized Lifetime Earnings for Oldest Relative to Youngest Workers 45

1.5 Percentage Changes Relative to Pre-Liberalization Steady State 46

2.1 Second-Stage Calibration 86

A.1 Calibrated Parameters 138

List of Figures

1.1 Menu of Technologies 30

1.2 Percentage Change in Average Income Relative to Pre-liberalization Along the Transition Path 35

1.3 Percentage Change in Average Income Relative to Pre-liberalization by Education Along the Transition Path 36

1.4 College Premium Along the Transition Path 41

1.5 Annualized Lifetime Earnings of Old Workers Relative to Young Workers . . 44

1.6 Education Cutoffs 48

1.7 Percentage Change in Skill Supply Along the Transition Path Relative to Pre-liberalization 49

2.1 Calibrated Prices 76

2.2 Empirical Age-Earnings Profile 79

2.3 Productivity and Average Income 81

2.4 Productivity by Experience and Age 85

2.5 Equilibrium Service Prices 88

2.6 Price Index and Relative Prices 89

2.7 Discounted Lifetime Earnings and Sectoral Entry 91

2.8 Increase in Average Age Attributable to Changes in World Prices 93

2.9 Change in Labor-Force-Participation Rate Attributatble to Changes in World Prices 95

2.10 Increase in Unemployment Rate Attributable to Changes in World Prices . . 96
2.11 Change in the Unemployment Rate Attributable to Changes in World Prices by Age 97
2.12 Average Income 99
2.13 Changes in Aggregate Welfare Attributable to Changes in World Prices . . . 101
2.14 Changes in Welfare Attributable to Changes in World Prices 102

Chapter 1

Skill Acquisition and the Dynamics of Trade Induced Inequality

This chapter 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.[1]

1.1 Introduction

The main objective of this chapter 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.[2] 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.[3]

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

[1] A version of this chapter was presented in conferences at the Federal Reserve Bank of Chicago and the University of Warwick.

[2] See 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.

[3] See Kwark and Rhee (1993) and Kim and Topel (1995).

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 acquisition. 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.[4]

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

[4]Other 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.

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 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.[5] 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.[6]

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.[7] In addition, Burstein and Vogel (2011) find that, of the increase in the skill premium associated with increased 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

[5] See 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.

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

[7] See, for example, Haltiwanger et al. (2004) and Wacziarg and Wallack (2004).

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 chapter 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 chapter contributes to a growing literature emphasizing the role of firm heterogeneity as a crucial link between trade and inequality.[8] 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 Kreickemeier (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 non-competitive wage setting is needed to rationalize inequality.[9]

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

[9] Verhoogen (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

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.[10]

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 chapter is also related to a nascent literature on the dynamic impact of trade on workers.[11] 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

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.

[10] Lileeva 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.

[11] There 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).

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) extend this model to incorporate sector-specific human capital as an added barrier to intersectoral mobility.[12] 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

[12] In 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.

versa.

1.2 The Model

1.2.1 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 (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.[13] 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.

1.2.2 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.[14]

In each period a firm must choose a technology, $v \in \{0, 1, ..., V\}$, 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 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

[13] Although 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 A.3.

[14] This is the Chaney (2007) variant of a Melitz model.

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.[15] 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.[16]

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

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

[15] As 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.

[16] An 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.

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.[17] The time subscript, t, is therefore suppressed in the firm-side analysis.

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,

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

[17] Because 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.

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

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

where

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

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.[18] 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 $\mu_{v(\varphi)} \leq \mu_{v(\varphi')}$.

[18] In 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.

Denote by $\tilde{v} \in \left\{0, 1, ..., \tilde{V}\right\}$ 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.[19] 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 $\Phi_{\tilde{v}} = \left(\varphi_{\tilde{v}}, \varphi_{\tilde{v}+1}\right)$, 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

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

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

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 $\tilde{A}_F = A_H \left(\tau w_\ell\right)^{1-\sigma}$. 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

[19] In general, not all the available technologies will be used in every period. In Appendix A.1, I provide an algorithm for determining which technologies are used in equilibrium and by which firms.

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,

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

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

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

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, $\left(\varphi_{\tilde{v}_x},\varphi_x\right)$, and exporting firms, $\left(\varphi_x,\varphi_{\tilde{v}_x+1}\right)$. 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 $\tilde{V}+1$.

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

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

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

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

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

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

Final-good supply is[20]

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

1.2.3 Workers

In this section, I develop an overlapping-generations model of endogenous skill acquisition. Workers are divided into age groups, $b \in \{1, 2, ..., B+1\}$. 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.[21] The measure of newborns is equal to the measure of deaths in each period so that there is no population growth.

[20] Because 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.

[21] As 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).

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, 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.[22]

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.[23] 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.

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

[23] In 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, $aw_{ht} > w_\ell$ 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.

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.[24] 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 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, $aw_{ht} > w_{\ell t}$ 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.

[24] In Appendix A.1 I provide an algorithm for this computation.

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, $\{W_t(a, b, e)\}_{t=s}^{\infty}$.[25]

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, respectively. 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,

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

Therefore, the supply of low-skill labor is

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

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,

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

[25] In Appendix A.1 I provide an algorithm for this computation.

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.

1.2.4 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,

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

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,

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

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

1.3 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.

1.3.1 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.[26]

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.[27]

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

[26] See Appendix A.2 for details on the size bins and CPS bins.

[27] Alternative 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.

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.

1.3.2 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 $\tau = 1.91$.[28]

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 off firms, M; and (4) the Pareto shape parameters for worker innate ability, θ_w, and firm

[28] This 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).

idiosyncratic productivity, θ_f.

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).[29] 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 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,

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

[29] The 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.

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.

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 A.2.

In order to calibrate the firm-level technologies, I assume that each of the thirty-two size bins described in Section 1.3.1 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.[30] 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, 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.[31]

[30] The 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 Trefler (2010) and Bustos (2007).

[31] Other choices of f_x in this interval do not influence the results.

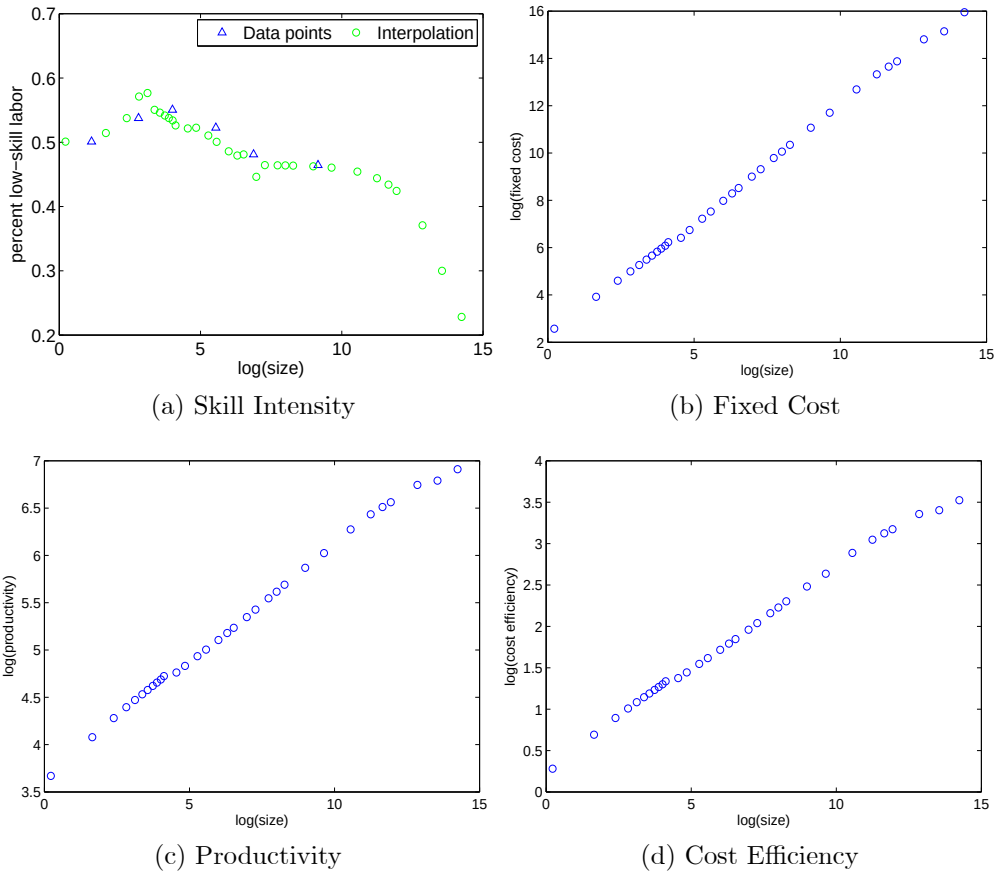
Results of Technology Calibration

The results of the technology calibration are summarized in Figure 1.1. 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 1.1a, 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 1.1a 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.[32] 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 1.1a instead of the interpolation used in the baseline calibration described in Appendix A.2. In both cases, the alternative interpolations have little quantitative impact on the results I report in Section 1.4.

[32] In 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 A.2.

Figure 1.1: Menu of Technologies



(a) Skill Intensity

(b) Fixed Cost

(c) Productivity

(d) Cost Efficiency

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 A.2.

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 1.1b and 1.1d. Thus, the model, and in particular the firm side of the model, is consistent with the dimensions of the data used in the calibration.

1.3.3 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.

1.4 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 A.3.

1.4.1 Results

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.[33] 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 1.2. 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.[34]

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

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

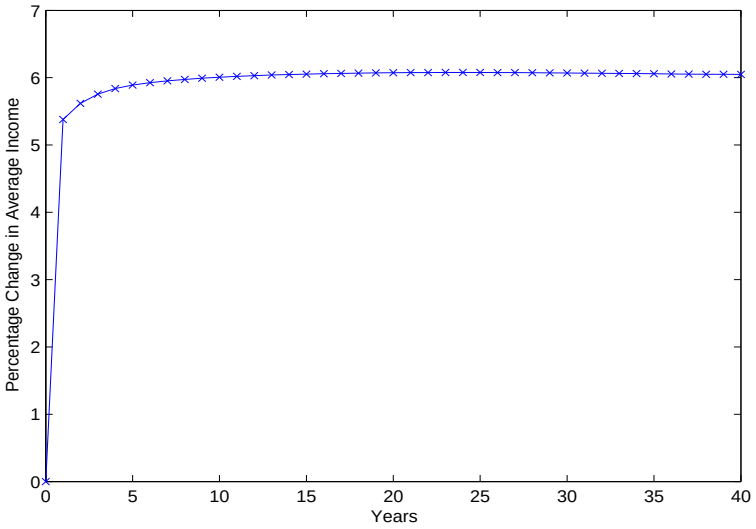
[34] This is not a general feature of the model. In particular, because of the inverted-U relationship between firm productivity and skill intensity, it is not necessarily the case that exporters are on average more skill-intensive than non-exporters. In general, this depends on the precise location 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.

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.[35]

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 2.12a 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.[36]

Figure 1.2: 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.

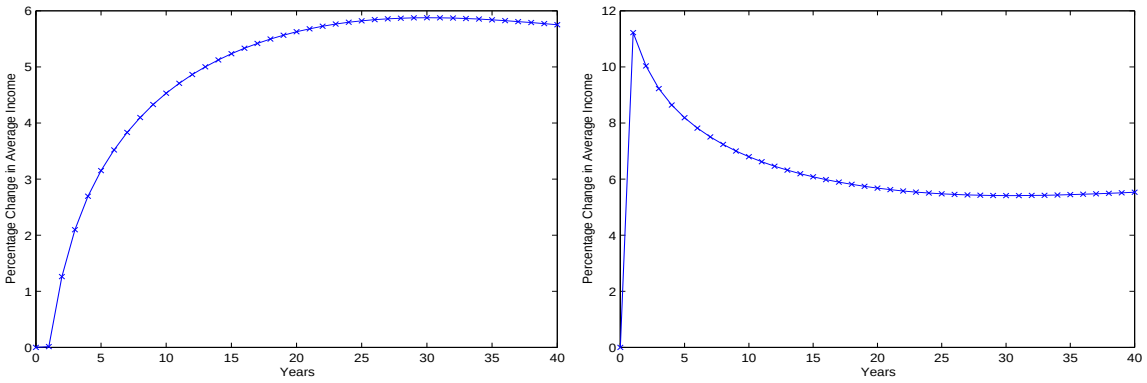
[35] Changes in the skill premium incentivize workers to pursue an education which increases the demand for teaching labor. Because the labor teaching requirement is more skill intensive than the aggregate skill intensity in the intermediate-good sector, increases in the number of students increase the 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.

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

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 1.3a and 1.3b 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 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.

Figure 1.3: Percentage Change in Average Income Relative to Pre-liberalization by Education Along the Transition Path



(a) Uneducated Workers (b) Educated Workers

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.

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.[37]

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 1.1 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,

[37] 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 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 last column shows that the gains are more uneven along the worker education dimension. Educated workers as a group gain 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 1.1: 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 1.2: 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 1.2 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 1.2 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.

Earnings Inequality

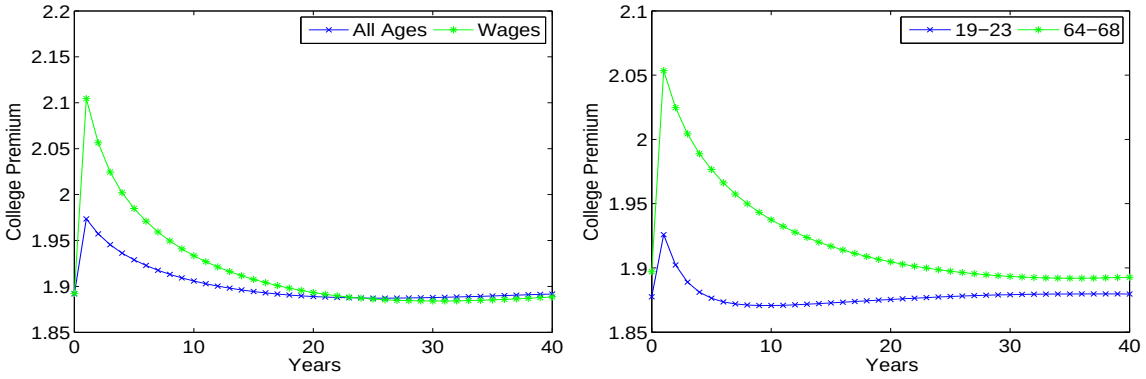
In this section, I consider the cross-sectional differentials in lifetime earnings among workers along the transition path. This is distinct from the differences in the gains from trade considered above because in this section I take into account the ex-post outcomes of workers' education decisions and 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.[38] 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.

[38] Care 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.

Inequality By Education

Figure 1.4 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.

Figure 1.4: College Premium Along the Transition Path



(a) College Premium vs. Income Premium

(b) College Premium Among Young and Old Workers

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.

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 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 1.3, the effect is reversed. Nevertheless, Figure 1.4b shows that earnings inequality is always higher among the older workers.

Table 1.3 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.

Table 1.3: 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.

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 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.

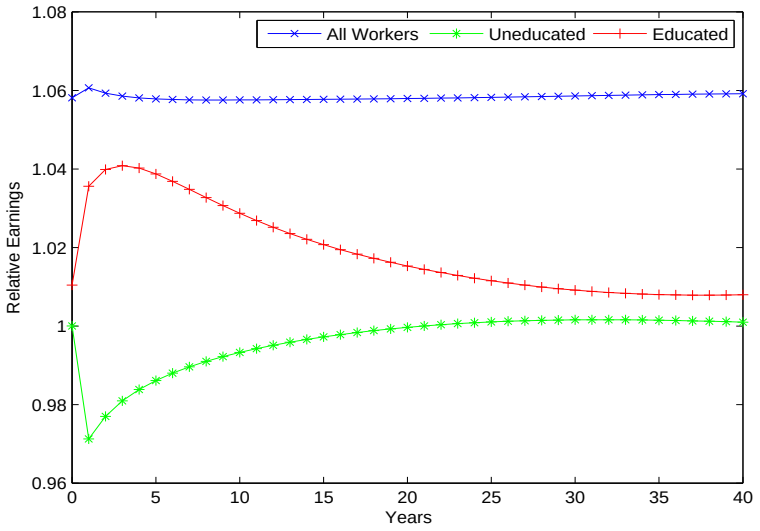
Finally, Figure 1.4a also plots the evolution of per-period income inequality, defined as the per-period average income of an educated worker relative to that of an uneducated worker. Although income inequality is often used as a proxy for lifetime-earnings inequality, a more meaningful measure of actual 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 1.5 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 1.5: 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 1.4 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 1.4: 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.

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 1.5 demonstrate the importance of the endogenous skill supply and the transition.[39]

Table 1.5: 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

[39]The 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.

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.

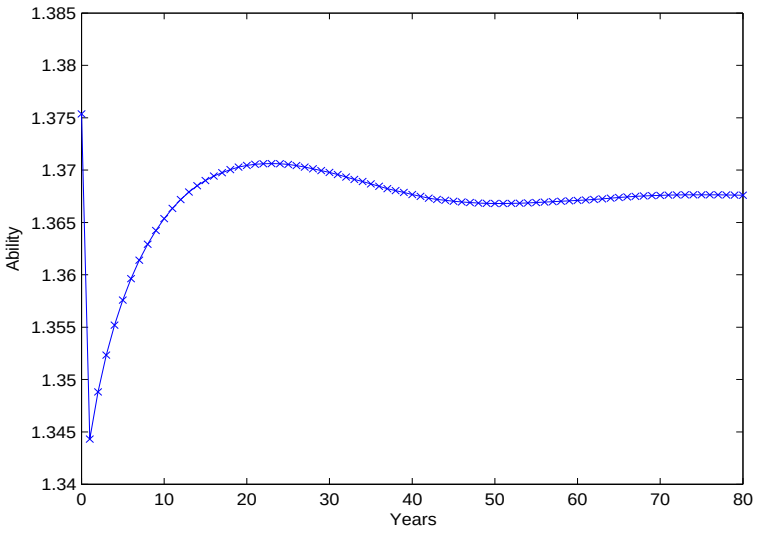
Finally, because all the dynamics in the model operate through changes in the skill supply, when the skill supply is fixed there are no dynamics. Therefore, ignoring either the transition or the endogeneity of the skill supply fails to account for the differential impact of trade across age groups. Old and young would be affected in the same way were there no transition.

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 1.6 shows the evolution of the education cutoffs for the youngest age group.

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 1.6 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

Figure 1.6: 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.

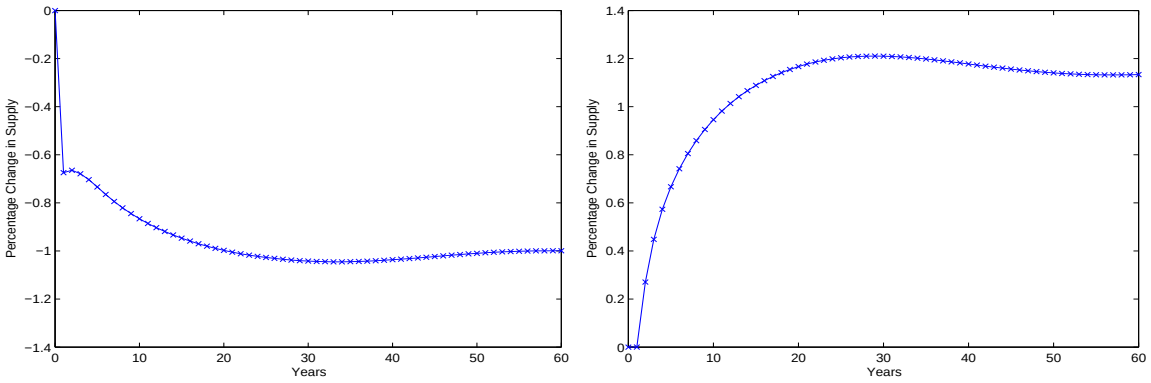
level, the education cutoff drops to below its long-run level. Thus, the cohort of young workers at the time of trade liberalization will be overeducated relative to the long run. As older undereducated workers die and the younger overeducated workers take their place, the scarcity of high-skill labor 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 1.7b 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.[40] These cycles explain the oscillations,

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

apparent in Figure 1.3, of the educated and uneducated workers' incomes.

Figure 1.7: Percentage Change in Skill Supply Along the Transition Path Relative to Pre-liberalization



(a) Low-Skill (b) High-Skill

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 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 1.7a.

1.5 Conclusion

The objective of this chapter 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 chapter 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.

Chapter 2

The Impact of World Price Changes on a Small Open Economy

This chapter studies the impact of changes in world prices of traded goods on a small open economy. I propose a dynamic model of sectoral choice and job search in which changing world prices affect sectoral entry patterns, unemployment and labor-force participation. I calibrate the model to Canadian data to quantify the impact of the changes in world prices of traded goods from 1993 to 2010. I find that the changes in prices over this time span increased unemployment in the manufacturing sector by 0.3 percentage points and decreased it in the agriculture and natural resource sector by over 1 percentage point. In addition, I find that workers in the nontraded service sector and the manufacturing sector lose from the changes. Since these workers represent approximately 95 percent of the Canadian workforce, the vast majority of Canadian workers have been hurt by changing world prices. In fact, aggregate welfare decreases by almost 7.5% in the short run and 5.5% in the long run due to the changes in world prices.

2.1 Introduction

Trade economists have long extolled the virtues of trade openness and the advantages of increased global integration. Indeed, much of the trade literature reads like a catalogue of the now countless mechanisms proposed for the gains from trade. Yet, despite the arguments in favor of globalization set forth by trade economists, ongoing global integration efforts continue to be met with skepticism and still inspire fear and anxiety among many workers. Clearly, the general public remains far from wholly convinced. Nor should they be convinced.

Most trade models, in order to illustrate the gains from trade, abstract away from many of the frictions that can potentially erode or even nullify the gains from trade. Although, the theoretical wage implications of globalization are relatively well understood, policy makers are concerned about other labor-market outcomes as well, such as unemployment and labor-force participation. Moreover, most trade models, by virtue of being frictionless, are implicitly long-run models. As a consequence, many trade models shed little light on the potential short- or medium-run implications of increased globalization. These shortcomings render many of the most widely used trade models ill-equipped to address many of the concerns raised in the public debate regarding global integration.

Indeed, recent findings lend credence to these concerns. Autor, Dorn and Hanson (2012) find that, at least for the United States, Chinese import penetration is associated with increases in unemployment rates and decreases in labor-force-participation rates. Furthermore, a recent literature has emphasized that a full rendering of the implications of global integration is impossible without accounting for transitional dynamics. This literature has identified sluggish labor-force reallocation across sectors as a key friction that can potentially erode the gains from trade in the short run. But, other frictions, for instance search frictions, may be important as well even in the long run.

The goal of this paper, therefore, is to take some steps in bridging the disconnect between the questions addressed by the existing trade literature and the concerns voiced in the public debate. In particular, I develop a dynamic overlapping-generations model of sectoral choice

and job search. I then embed this model in a general-equilibrium framework that allows me to study the implications of global changes on a small open economy with a nontraded sector both in the short run and the long run.

In the model, workers, at birth, decide which sector to enter by comparing the expected-discounted-lifetime utility afforded by each sector to their idiosyncratic cost of entry into each sector. Once within the sector, workers obtain employment by searching for a job, which involves drawing a job opportunity at a utility cost. All jobs within a sector offer the same wage conditional on the worker's age and experience, but differ in the utility cost they entail. As a result, workers may reject some job opportunities in the hopes of finding a better job in the future. In addition, if the expected value of searching does not outweigh the cost of search, a worker may prefer to drop out of the work force altogether rather than incur the search cost. Thus, in addition to sectoral choice, workers face a job-opportunity-acceptance decision as well as a labor-force-participation decision.

How do changes in the global economy, modeled as changes in the world prices of traded goods, affect the small open economy? Consider a sector which experiences a price decline. The direct effect is that it lowers wages in that sector. This renders some jobs, previously acceptable, unacceptable. As a result, each search episode is less likely to result in employment, and some workers may choose to drop out of the labor force altogether. In addition, if the sector becomes, as a result of the price decline, less attractive than other sectors, the sector will shrink as fewer young workers choose to enter it. This will affect the age composition of the work force, and as I show, lead to nontrivial dynamics in the average income earned within each sector.

The effect of changes in the global economy, however, does not end with the changes in the relative prices of traded goods. Due to its effect on workers' incomes and sectoral entry patterns, changes in world prices affect both the demand for and the supply of the nontraded good. This general-equilibrium effect impacts workers in all sector. This is because ultimately what matters to workers is the price of the good they produce relative to the price

of the good they consume. The impact of these general-equilibrium effects depends heavily on the parametrization of the model, and it is therefore difficult to know their qualitative impact in general even in the long run, much less along the economy's transition path. Thus, although I derive some of the model's properties analytically, the analytic results serve mainly to deepen our intuition for the mechanisms embedded within the model. Ultimately they underscore the importance both of a rigorous quantitative analysis and of a need to account for transitional dynamics as shifts in sectoral entry patterns move the real prices of all goods, traded and nontraded.

To calibrate the model, I divide the Canadian economy into three traded sectors, agriculture, natural resources and manufacturing, and one nontraded sector, services. Parameters are chosen to match important features of the Canadian economy in each of these four sectors. I use the calibrated model to quantify the impact of changes in world prices from 1993 until 2010 on the Canadian economy and Canadian workers. I do so by simulating the calibrated model first with the actual time series of prices in the traded sectors and then with prices in the traded sectors counterfactually fixed at their 1992 levels. The counterfactual simulation is interpreted as what would have happened to the Canadian economy but for the changes in world prices. The difference in outcomes between the two simulations is therefore interpreted as the impact of changes in world prices on the Canadian economy.

Since I do not model the global general equilibrium that determines world prices in the traded sectors, I cannot attribute my findings to any specific cause. Nevertheless, the price changes that I derive are consistent with how we might expect the growth of China and the rest of the developing world to affect world prices. I find that manufacturing prices declined by approximately 20 percent and natural resource prices increased by nearly 50 percent between 1993 and 2010. Growth in the developing world, which led to increased energy demands as well as increased supply of manufactures to the rest of the world, is a plausible explanation for both of those price changes. Thus, although the model does not

identify them as such, it is reasonable to interpret the results as, at least in part, the effect of growth in the developing world on Canadian workers.

I find that the world price changes have no appreciable effect on aggregate unemployment in Canada. However, when unemployment is decomposed by sector the story changes. Unemployment in the manufacturing and service sectors increases by approximately 0.3 percentage points and 0.1 percentage points in the long run, while it decreases by over 1 percentage point in both the natural resource and agriculture sectors. Labor-force participation increases substantially in the natural resource sector, but is almost completely unaffected in the other sectors.

I find that approximately a 5 percentage point decline in the share of workers entering the manufacturing sector is attributable to the 20 percent decline in the price of manufactures. These workers instead enter the other three sectors, but predominantly the service sector. The change in the entry pattern leads to a decrease in the price of services, and ultimately the change in world prices is responsible for a 15 percent decline in the price of services. As a consequence, even though the changes in world prices decreases real prices in the economy by 10 percent, both manufacturing and service sector workers are hurt by the change in world prices.

In the long run, the average worker in the natural resource sector is approximately 8 percent better off as compared to the average worker in the natural resource sector had prices remained unchanged. However, the average worker in the agriculture, services and manufacturing sectors are, respectively, 3, 8 and 9 percent worse off. Importantly, however, because of the distribution of idiosyncratic sector-entry costs, workers do not simply enter the natural resource sector en mass despite the increased attractiveness of that sector. Instead, most workers continue to enter the manufacturing and service sectors despite the decrease in those sectors' attractiveness. As a result, an aggregate welfare decline of 7.5% in the short run and 5.5% in the long run is attributable to changes in world prices.

Relation to the Literature

The research in this chapter is most directly inspired by the aforementioned findings of Autor, Dorn and Hanson (2012). Their focus is on manufacturing workers in the United States, whereas I focus on Canadian workers in all sectors. However, methodologically the key difference is that their findings are based on reduced-form estimation, while I employ a structural analysis. This allows me to identify the full impact of changes in the global economy rather than just the differential impact across geographical regions as in Autor, Dorn and Hanson (2012). In addition, because I follow the Canadian economy from 1993 and on, I am able to quantify the impact of global changes both in the short run and the long run, rather than just at one particular point in time, which, may, as my results show, be misleading with respect to the overall impact of trade shocks.

This chapter is related to a strand of the literature which emphasizes the importance of accounting for endogenous sectoral choice in an economy's adjustment to trade shocks. In early work, Matsuyama (1992) studies a model in which adjustment to trade shocks occurs through the process of demographic change. More recently, 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) extend this model to incorporate sector-specific human capital as an added barrier to intersectoral mobility.

A recent literature has argued that incorporating labor-market frictions into models of trade provides important insights into the impact of trade shocks on an economy. One strand of this literature, starting with Davidson and Matusz (1999) and more recently Helpman and Itskhoki (2011) among others, has studied how labor-market tightness responds endogenously to trade shocks. In another strand of this literature, Davidson and Matusz (2000) and Cosar (2010) highlight the role of labor-market frictions in explaining the sluggishness of an economy's adjustment to trade liberalization. This strand of the literature emphasizes that labor-market frictions represent adjustment costs which can substantially undermine

the benefits of trade liberalization. My model, emphasizes both the adjustment aspects of unemployment, as well the long run impacts.

Although aspects of my model exist in the literature, the key innovation is to incorporate all of these elements in a tractable dynamic model. In particular, this is the first dynamic model in the trade literature that endogenizes sectoral choice, unemployment and labor-force-participation. My novel identification methodology including calibration of the model to match the Canadian economy from 1993-2010 allows me to identify the impact of global changes on the Canadian economy over time. Furthermore, my methodology does not rely, as the existing literature does, on the unrealistic assumption that the economy is in a steady state at the time of the trade shock or that other changes, in particular productivity changes, do not occur along the economy's transition path.

2.2 Model

2.2.1 Setup

Consider a model of a small open economy set in discrete time with periods indexed by t. The economy consists of S sectors indexed by $s \in \{1, ..., S\}$, where sectors 1 through $S-1$ are traded and sector S is nontraded. Goods differ across sectors, but are homogeneous within a sector. The sector-level goods are aggregated into a final nontraded good according to a constant-elasticity-of-substitution (CES) aggregator. In any period t, the supply of the final good in the economy is

$$Q_t = \left[\sum_{s \in S} \alpha_s^{1/\sigma} q_{st}^{(\sigma-1)/\sigma} \right]^{\sigma/(\sigma-1)},$$

where q_{st} is the quantity supplied of the good produced is sector s in period t, σ is the elasticity of substitution between the goods produced in the various sectors and α_s is the relative importance of the sector s good in the final good, with $\sum_{s=1}^{S} \alpha_s = 1$. Given this

CES structure, the price of the final good in period t is

$$P_t = \left[\sum_{s \in S} \alpha_s {p_{st}}^{1-\sigma} \right]^{1/(1-\sigma)},$$

where p_{st} is the price of the sector s homogeneous good in period t. Because the economy is a small open one, the prices p_{st} are exogenously given by world prices for $s < S$, but p_{St} is endogenously determined by supply of and demand for the nontraded good within the economy.

Labor is the only factor of production, and it is supplied by discounted-expected-utility-maximizing workers. Workers, who are the only type of agent in the model, derive utility from consumption of the final good, but suffer disutility from work-related activities that include both work and search to be described below. Workers discount utility by a factor β, and their periodic utility is equal to their consumption of the final good, or their real income, less their disutility from work.[1] Finally, workers are price takers and have perfect foresight.

2.2.2 Workers

In this section, I develop an overlapping-generations model of sectoral choice and job search. This model will serve as the production technology to be embedded in the general-equilibrium setup just described. Workers are born into age group $a = 1$ and live for A periods. Upon birth, each worker draws S costs, $\{C_s\}_{s=1}^{S}$, where C_s represents the utility cost of entering sector s. The cost draws are independently and normally, though not identically, distributed, so that $C_s \sim N(\mu_s, \theta_s)$. The population is normalized to unity, and there is no population growth, so that in each period $1/A$ workers are born while the same measure of workers die.

The first decision faced by a worker is which sector to enter. Once the initial sectoral choice has been made, sector switching is not allowed.[2] Thus, a worker may search for a job

[1] This formulation implies that all disutility costs are in terms of the final good.

[2] See Matsuyama (1992).

and work only in the sector they choose at birth.[3] Workers choose the sector that affords them the highest discounted expected utility net of their idiosyncratic cost of entering the sector.[4] In order to determine the value of entering a sector, we now turn to the details of the production technology within each sector.

Labor Force Participation and Search

In this subsection, I consider workers within a given sector after their initial sectoral choice. Thus, in order to avoid clutter, for the remainder of this subsection, I drop the subscript s from the development of the model, with the understanding that the parameters introduced are sector specific. Each worker is characterized by a triplet, (a, e, j), where a is the worker's age, e is the worker's experience measured in periods, and j indicates the type of job opportunity the worker currently holds. Specifically, j indicates the job-specific disutility incurred per period while working at job j. By convention, $j = \infty$ for a worker with no current job opportunity, which can occur either if the worker is unemployed or if the worker is not participating in the labor force. Since surely no worker would accept a job opportunity that affords her infinite disutility and no worker is ever forced to accept a job opportunity, holding a job opportunity $j = \infty$ is equivalent from the point of view of the worker to having no job opportunity at all. I will refer to such a worker as jobless. Every worker, immediately upon entry into the sector, is characterized by the triplet $(1, 0, \infty)$. In other words, fresh entrants into a sector are in the youngest age group, have no work experience and are jobless.

Consider a jobless worker in period t characterized by the triplet (a, e, ∞). Such a worker can choose to do nothing, i.e. drop out of the labor force. In such a case, the worker's period t utility is zero, and she enters the following period characterized by the triplet $(a + 1, e, \infty)$. Unlike sectoral choice, dropping out of the labor force is reversible, and there is no difference

[3] Wacziarg and Wallack (2004) find that trade liberalization leads to little or no worker reallocations across industries.

[4] Since workers cannot switch sectors the cost of entering a sector can alternatively be interpreted as the discounted present value of a utility cost that the worker incurs every period in the sector.

between jobless workers who were in the labor force in the prior period and those who were not, modulo age and experience.

Alternatively, a jobless worker can choose to search for a job, in which case the worker is in the labor force and therefore unemployed. Search entails drawing a job opportunity, j, from a job-opportunity distribution at a utility cost of c. The job opportunity that is drawn is available to the worker in the subsequent period. Thus, if a worker chooses to draw a job opportunity, which I will refer to interchangeably as searching, her periodic utility is $-c$. Workers will potentially be willing to incur this temporary utility hit because of the prospects of finding a desirable job for future periods.

Search in this model does not involve matching searching workers with firms, rather workers can be thought of as self-employed, so that they own anything they produce. For a given age, a, and experience, e, all jobs are equally productive. However, jobs are not identical because they differ in the utility cost incurred by the worker while employed at the job, as described above. Specifically, in each period that a worker is employed at job j, she incurs a disutility of $c + j$. The job-specific disutility, j, is assumed to be normally distributed with a zero mean and a standard deviation of η, or $j \sim N(0, \eta)$. Since the worker foregoes leisure while she is searching or working, c can be interpreted as the utility cost of foregone leisure. The job-specific disutility, j, differs across jobs because jobs potentially differ in their commute time, general unpleasantness (or pleasantness) and any number of other dimensions which directly impact the utility of the worker.[5] If a jobless worker decides to search and draws a job opportunity j, she enters the subsequent period characterized by the triplet $(a + 1, e, j)$.

Finally, consider a worker with a job opportunity $j < \infty$, characterized by the triplet (a, e, j). The worker can either reject or accept the job opportunity. If the worker rejects the job opportunity, she immediately re-enters the jobless worker pool. She is then characterized by the triplet (a, e, ∞) and is allowed to search, if she so desires, within that same period. In

[5] The normal distribution can be justified by appeal to the Central Limit Theorem on the grounds that the disutility incurred by the worker is a sum of a large number of independently drawn disutilities.

other words, she is identical to all the other workers in her age group and with her experience who began the period jobless. Thus, job-opportunity rejection can be thought of as occurring between two periods, or alternatively, since workers have perfect foresight, immediately upon drawing the job opportunity.

If instead the worker accepts the job opportunity, she receives, at a utility cost of $c + j$, an income of $p_t A_t h(a, e)$, where $h(a, e)$ is the human capital function of the worker that translates age and experience into worker productivity, p_t is the period t price of the good produced in the sector under consideration and A_t is the period t productivity in the sector. The function $h(a, e)$ is equal to $g(a) k(e)$, where $g(\cdot)$ is a weakly decreasing function of age, and $k(\cdot)$ is a weakly increasing function of experience. Thus, a worker employed at job j receives period t utility equal to $p_t A_t h(a, e) - c - j$. The worker gains experience while on the job, but with probability λ she experiences an exogenous job-separation shock and loses her job in the subsequent period, i.e. she does not have a job opportunity with which to begin the subsequent period. Thus, with probability λ the worker enters the next period characterized by the triplet $(a + 1, e + 1, \infty)$ and with the complementary probability enters the next period characterized by $(a + 1, e + 1, j)$. In this model, job separations can happen either exogenously, due to a job-separation shock, or endogenously if a worker rejects the job she held in the previous period.

To summarize the mechanics of the job-search model, workers can either hold a job opportunity or be jobless. A worker may be jobless for one of four reasons. The first reason is that the worker is a newborn and has yet to search for a job. The second reason is that the worker had a job opportunity and rejected it. The third reason is that the worker had a job and lost it exogenously. The final reason is that the worker was jobless in the previous period and decided not to search. A worker may hold a job opportunity for one of two reasons. The first reason is that the worker was jobless in the previous period and drew a job opportunity. The second reason is that the worker had a job in the previous period and did not experience a job-separation shock, so that the job opportunity carried over to the

current period.

Worker Optimization

In this subsection, I consider the workers' optimization problem. In solving this problem, the worker takes as given the path of prices and productivities in each sector during her lifetime, $\left\{\{p_{st}\}_{s=1}^{S}\right\}_{t=b}^{b+A}$ and $\left\{\{A_{st}\}_{s=1}^{S}\right\}_{t=b}^{b+A}$, where b is the birth period of the worker. The worker's problem can be solved in two stages. First, the worker finds the optimal search policy *given* a sectoral choice. The optimal search policy consists of, for each age-experience pair, a decision rule for jobless workers whether or not to draw a job opportunity and a decision rule for workers with a job opportunity whether or not to accept it. Second, given the within-sector optimization, the worker chooses the sector that offers the maximal discounted expected utility. In Appendix B.2.1, I provide an algorithm for computing worker values and policies.

Consider a worker holding a job opportunity with associated utility cost j. For a given age-experience pair, (a, e), the lower is j the more attractive is the job opportunity. Thus, the job acceptance decision rule exhibits a cutoff property such that in period t in sector s all jobs above the cutoff $j_{ts}(a, e)$ are rejected and those below it are accepted.

A job is valuable to a worker both for its income and as an investment. It provides value as an investment for two reasons. First, it increases the worker's experience and, as a consequence, her future productivity. Second, it provides the worker an option value in the form of a job opportunity in the subsequent period. Since younger workers have a longer horizon ahead of them, the investment value of a job is greater for younger workers. Furthermore, since productivity is weakly decreasing in age, for a given experience level, the income value of a job is also greater for younger workers. Therefore, conditional on experience the benefit of a job is decreasing with age. As a consequence, the value of searching is also decreasing with age. However, the cost of search is the same for all workers independent of age. Thus, there exists an age cutoff for a given experience level e, such that in period t in

sector s all jobless workers above the age cutoff $a_{ts}(e)$ choose to drop out of the labor force and all those below the cutoff choose to search for a job.

Finally, consider the sectoral choice of a worker. Every new entrant into a sector is identical, and therefore, the discounted expected utility of entry into a sector is equal for all workers. Denote the value of entry into the S sectors for a worker born in period t by $\{V_{ts}\}_{s=1}^{S}$. Thus, given idiosyncratic sector-entry cost draws, $\{C_s\}_{s=1}^{S}$, a worker chooses to enter the sector s^* that satisfies

$$s^* = \arg\max_s \{V_{ts} - C_s\}.$$

Not surprisingly then, an increase in the value of entering one sector relative to the others will result in an increase in the number of entrants into that sector. However, because the idiosyncratic cost draws are normally distributed, for any set $\{V_{ts}\}_{s=1}^{S}$, there will always be entrants into every sector.[6] This is consistent with workers entering every sector even in the event that some sectors offer vastly higher discounted expected utilities than other sectors.

Worker Distribution

The distribution of workers in any period t is a four-dimensional object $W_t(a, e, j, s)$, which indicates how many workers there are of age a with experience e holding job opportunity j in sector s. An initial worker distribution $W_{t_0}(a, e, j, s)$ and worker policy are sufficient to compute the entire future path of worker distributions, $\{W_{t_0}(a, e, j, s)\}_{t=t_0}^{\infty}$. I describe an algorithm for this computation in Appendix B.2.2.

Knowledge of these worker distributions allows me to compute many of the outcomes that I am particularly interested in studying in this paper, including the labor-force-participation rate, the unemployment rate and average income by age, experience and sector. However, the only reason that knowledge of the distribution of workers is necessary to compute an

[6] Of course, this would be true for any unbounded distribution.

equilibrium is in order to compute supply of and demand for the nontraded good. The quantity of good s produced in the economy is

$$q_{ts} = A_{ts} \sum_{a=1}^{A} \sum_{e=1}^{A} h\left(a,e\right) \int_{-\infty}^{j_s(a,e)} W_t\left(a,e,j,s\right) dj,$$

where the integral sums over all the workers with accepted job opportunities for age a and experience e. This integral is multiplied by the human capital function $h\left(a,e\right)$ because the quantity produced by these workers depends on their age and experience.

Supply of the nontraded good in period t, q_{tS},depends only on sector S. However, demand for the nontraded good in period t is determined by total income in the economy, I_t, summed over all sectors. Total income in the economy is

$$I_t = \sum_{s=1}^{S} p_{ts} q_{ts}.$$

To summarize, given a path of productivities and prices in each sector, $\left\{\left\{A_{ts}\right\}_{s=1}^{S}\right\}_{t=t_0}^{\infty}$ and $\left\{\left\{p_{ts}\right\}_{s=1}^{S}\right\}_{t=t_0}^{\infty}$, worker optimization implies a path of worker policies in the form of decision rules regarding sectoral entry, labor-force participation and job-opportunity acceptance. This path of worker decision rules, given an initial worker distribution, yields the entire future path of worker distributions. Finally, these worker distributions together with worker decision rules regarding job-opportunity acceptance yield both total income in the economy and the supply of the nontraded good.

2.2.3 Equilibrium

An equilibrium in the open economy is defined for an initial worker distribution, W_{t_0}, a path of productivities in each sector, $\left\{\left\{A_{ts}\right\}_{s=1}^{S}\right\}_{t=t_0}^{\infty}$, and world prices in the traded sectors $\left\{\left\{p_{ts}\right\}_{s=1}^{S-1}\right\}_{t=t_0}^{\infty}$. Given this initial worker distribution and paths for productivities and world

prices, an equilibrium is fully characterized by a path of prices for the nontraded good, $\{p_{tS}\}_{t=t_0}^{\infty}$. As described in the previous sections, the exogenously given path of productivities and world prices together with the path of nontraded sector prices are sufficient to compute optimal worker policy in every period. The path of nontraded good prices are an equilibrium if, given optimal worker policy implied by this path, the nontraded good market clears in every period.

Because the final good is a CES aggregate of the sectoral goods, the demand for the nontraded good, given income I_t, is

$$D_{tS} = \frac{I_t \alpha_S p_S^{-\sigma}}{P^{1-\sigma}}.$$

Thus, equilibrium requires that $D_{tS} = q_{tS}$. A steady-state equilibrium of the small open economy is one in which all the endogenous variables are constant over time.

2.3 Mechanisms

Before moving on to the calibration and the quantitative analysis, it is worth discussing, through the lens of the theoretical model, the effect of changes in the global economy on worker decisions and outcomes. Changes in the global economy are modeled as changes in the relative prices of the traded goods, which occur through an unmodeled global general equilibrium. The small open economy assumption, however, ensures that changes in the economy under consideration have no bearing on the global general equilibrium. This is the sense in which world prices in the traded sectors are exogenous from the point of view of the small open economy. For the remainder of this section, I normalize the price of the final good to unity, so that all sector level prices should be interpreted as real prices, that is the price of the sectoral good relative to the price of the final good.

2.3.1 Unemployment and Labor-Market Participation

How might we expect changes in the global economy to affect worker decisions? Suppose that changes in the global economy, lead to a decrease in the relative price of a certain good. Because this decreases the value of the goods produced by the workers in this sector, the price change will have a direct deleterious impact on worker real incomes even though the change has no effect on the quantity produced by these workers.

The decrease in the income offered by jobs in the sector experiencing a decreased price makes jobs in this sector less attractive. For jobless workers, this decreases the value of search and increases the likelihood of a worker choosing to drop out of the labor force altogether rather than incurring the search cost. The lower incomes also affect the job opportunities that are accepted. This is because when jobs offer less pecuniary compensation, workers are no longer willing to incur the same utility cost of working as they would with higher compensation. The intuition just outlined is formalized in the following proposition.[7]

Proposition 2.3.1 *Consider two paths of prices and productivities in sector s, indexed by* 1 *and* 2*, where $A_{1ts}p_{1ts} > A_{2ts}p_{2ts}$ for all t. The following statements hold:*

1. *Conditional on experience e, the labor-force-participation age cutoff, $a_{ts}(e)$, is greater for path* 1 *than for path* 2 *for all t.*
2. *Conditional on age, a, and experience, e, the job-opportunity acceptance cutoff, $j_{ts}(a, e)$, is greater for path* 1 *than for path* 2 *for all t.*

The impact of changes in world prices and the associated changes in worker decisions will have an impact on both the unemployment and labor force participation rates in the sector that has undergone a change in the price of its good. Returning to the case of a relative-price decrease, this decrease lowers the labor-force-participation age cutoff, as shown in Proposition 2.3.1. Since fewer jobless workers choose to search, the result is a decrease in

[7] Although I will be mostly interested in the effect of price changes, I state the propositions in this section in terms of changes in $p_s A_s$. This is because from the worker's perspective, it is this product that matters. Whether the workers earnings change because of a change in productivity or whether it is because of a change in price is immaterial from the individual worker's point of view.

the labor-force-participation rate. Furthermore, as shown in Proposition 2.3.1, this price decrease will lower the job-opportunity acceptance rates among searchers. Thus, any given period of search is less likely to end in employment. This will lower employment relative to the size of the population within a sector. These conclusions are summarized in the following proposition:

Proposition 2.3.2 *Consider two steady states in sector s, indexed by* 1 *and* 2. *If* $A_s^1 p_s^1 > A_s^2 p_s^2$, *then:*

1. *The labor-force-participation rate is higher in steady state* 1 *than in steady state* 2.
2. *The employment-to-population ratio is higher in steady state* 1 *than in steady state* 2.

The model does not yield a sharp prediction for the change in the unemployment rate or average income across steady states. Consider first the unemployment rate when the relative price decreases. On the one hand, each searcher is less likely to accept a job opportunity she draws, which will push the unemployment rate up. On the other hand, fewer of the jobless will participate in the labor force, which will push the unemployment rate down. Without more parameter restrictions, the model is silent as to which effect dominates. Both of these effects, however, push the ratio of jobs to population down, which is precisely part 2 of the proposition.

Similarly, two effects govern the average income in a sector. On the one hand, an increase in the price of the good will increase the income of workers conditional on experience. On the other hand, the higher incomes might induce increased job acceptance and employment among workers with low human capital, $h(a, e)$. This effect will tend to decrease average income in the sector. As in the case of the unemployment rate, the dominant effect will depend on the specific parameterization of the model.

The reason that Proposition 2.3.2 is stated only for the case of steady-state equilibria is that the effect of a price change on unemployment and labor-force-participation rates depends on the prevailing worker distribution prior to the price change. Nevertheless, the

basic intuition underlying Proposition 2.3.2 holds true even when the economy is not in a steady state. However, even if we can pin down the initial worker distribution by considering the transition path from one steady state to another, Proposition 2.3.2 will not necessarily hold along the transition path for two reasons, compositional effects and general equilibrium effects. I discuss the impact of these two effects in the following subsections.

2.3.2 Compositional Effects

Consider an economy in steady state. Within each sector each age group is equally represented, i.e. $1/A$ proportion of workers in each sector belong to each age group. Now suppose that a once-and-for-all change in the global economy lowers the relative price of a good in one of the sectors. This will diminish the attractiveness of entry into the negatively impacted sector, and, as a consequence, fewer young workers will enter this sector than in the steady state prevailing prior to the price change. This will increase the proportion of older workers in the sector and decrease the proportion of younger workers. As a result, during the transition, older workers will be overrepresented in sectors hit by a negative shock and underrepresented in sectors experiencing a positive shock.

Consider now how these compositional changes affect average sector-level outcomes in a negatively impacted sector. Take, for example, average earnings in a sector. If older workers earn higher incomes than younger workers, because they, on average, have more years of work experience, then the compositional effect will lead to a higher average income relative to the average income that will prevail in the sector in the new steady state. In fact, if the compositional effect is sufficiently strong, it is possible that during the economy's transition to the new steady state average income will actually surpass its preshock level, and average income will exhibit an inverted-U shape pattern through the transition path. The same logic implies that the compositional effect would have the opposite effect on a positively impacted sector. The previous analysis is by no means limited to the case of average earnings. The same intuition can be brought to bear on any outcome whose average differs across age

groups and is therefore subject to these compositional effects. In particular, this intuition will be useful in explaining the dynamics of labor-force participation and unemployment in the quantitative analysis.

This discussion highlights the importance of accounting for the economy's transition path and the associated compositional changes along it. A decrease in the average income of a sector may incorrectly be viewed as a sign that workers in that sector are worse off. However, if this decrease in the average income is only a result of a compositional effect, then the decrease is inconsequential from the individual worker's point of view. The structure of the model will allow me to decompose the change in sector-level outcomes into changes resulting directly from changes in global prices, and those resulting only indirectly through the compositional changes within a sector.

2.3.3 General Equilibrium Effects

Until now the discussion has focused on the effect of price changes within a specific sector, while abstracting from the effects outside of the sector. In particular, Proposition 2.3.2 takes the steady-state price in a sector as given, even though these prices will, in general, be a result of the general equilibrium effects within the economy. In the absence of a nontraded sector, these general equilibrium effects are nonexistent and relative prices are determined by forces outside of the small open economy. In fact, if all sectors are traded, then a modified version of Proposition 2.3.2 holds.

Proposition 2.3.3 *Consider an economy in steady state in which all sectors are traded. Now suppose that there is a change in relative world prices. In all future periods, conditional on age, relative to the old steady state:*

1. *The labor-force-participation rate will increase (decrease) in sectors that experience a price increase (decrease).*
2. *The employment-to-population ratio will be increase (decrease) in sectors that experience a price increase (decrease).*

Because of the compositional effects discussed above, along the transition path Proposition 2.3.3 must be stated in terms of within-age-group averages rather than averages covering the entire sector as in Proposition 2.3.2. Understanding why this proposition fails when one of the sectors is nontraded gets to the heart of the importance of the general-equilibrium effects in the model. What matters to an individual worker is the price of the good she produces relative to the price of the final good. In a small open economy, if all sectors are traded, nothing that happens in the economy affects this relative price. It is all determined in the global general equilibrium independently of anything that happens in the small open economy. With a nontraded sector this is no longer the case.

When relative world prices change in the traded sectors, this affects both the supply of and the demand for the nontraded-sector good. Demand is directly affected because the changes in world prices affect income in the economy. Whether demand will increase or decrease as a result depends on many factors, including the relative sizes of traded sectors. If the relative price increased in a relatively large sector and decreased in a relatively small sector, this would tend to increase the income in the economy, thereby increasing demand for the nontraded good, and vice versa if the price changes were reversed. Thus, the effect of a change in world prices on demand for the locally-produced nontraded good is ambiguous.

Supply of the nontraded good is affected by changes in relative world prices through two channels. The first channel is that changes in world prices affect the price of the nontraded good relative to the price of the final good at the prevailing nontraded good price. If the relative world price increases in a sector which is relatively important in the CES aggregator, i.e. high α_s, and decreased in a low α_s sector, this will tend to make the final good more expensive relative to the nontraded good. This will make work less attractive in the nontraded sector, as per Proposition 2.3.1, thereby decreasing supply of the nontraded sector good. If the relative sizes of the α's were reversed the effect too would be reversed, making this effect ambiguous in terms of its impact on supply of the nontraded good.

The second channel through which a change in world prices affects supply of the nontraded good is through its effect on workers' sectoral choice. Depending on the parameterization of the distribution of sector-entry cost draws and depending on which traded sectors experience relative price increases or decreases, workers may be more or less likely to enter the nontraded sector following a change in world prices. If, for example, fewer workers enter the nontraded good sector, its supply will decrease, thereby leading to a price increase.

The sum total of all these effects is that not only does a change in world prices have an ambiguous effect on the price of the nontraded good, but it also has an ambiguous effect on the price of a traded good relative to the price of the final good. It is possible that a traded good whose relative world price increases actually experiences a price decrease relative to the final good if the price of the nontraded good increases sufficiently. Thus, when relative world prices change, the model, without further restrictions, does not provide a clear prediction regarding the change in prices relative to the price of the final good, even in steady state. This is the reason that Proposition 2.3.3 fails in the presence of a nontraded sector. This is before we even consider the complicated general equilibrium interactions for an economy along a transition path outside of steady state. The quantitative analysis, with a calibrated model, is therefore necessary to understand the general-equilibrium impact of changes in world prices.

2.4 Methodology

The goal of the quantitative analysis is to quantify the impact of changes in the global economy on a small open economy. Specifically, in this chapter, I focus on the effect of global changes on the Canadian economy and Canadian workers. The premise of my analysis is that events in Canada do not affect world prices in traded sectors. Instead these prices are determined in a global general equilibrium which remains unmodeled here, but the precise details of which, beyond the prices they ultimately determine, have no bearing on Canadian

workers. The analysis, therefore, hinges critically on identifying the evolution of the prices of traded-sector goods.

In order to identify the impact of global changes on the Canadian economy, I first solve the model with the actual time series of traded-sector prices, $\left\{\{p_{ts}\}_{s=1}^{3}\right\}_{t=1993}^{\infty}$, and productivities in all four sectors, $\left\{\{A_{ts}\}_{s=1}^{4}\right\}_{t=1993}^{\infty}$. This simulation, which I will call the actual simulation or actual equilibrium, involves finding a path of prices for the service-sector good, $\{p_{tS}\}_{t=1993}^{\infty}$, that clears the market for services in each period. I then solve the model with traded-sector prices counterfactually fixed at their 1992 levels. Because the nontraded-sector-good price is determined endogenously, in the counterfactual simulation these prices do not stay fixed at 1992 levels, unlike the prices of the other goods. Thus, this simulation, which I will call the counterfactual simulation or counterfactual equilibrium, involves finding another path of service sector prices. In both simulations workers are assumed to have perfect foresight regarding the evolution of prices after 1992, and in both simulations I solve the for the economy's entire dynamic transition path until a terminal steady state is reached.

All the model's parameters are held fixed across the two simulations, except for the world prices in the traded sectors. Since I assume that Canada is a small open economy, I interpret the outcomes in the counterfactual simulation as what would have happened in the Canadian economy but for changes in world prices of traded goods. As a result, the difference in outcomes between the actual and counterfactual simulations are identified as the effect of global changes on the Canadian economy.

Before I begin mapping the theoretical model to the data, I must first specify what is meant by a sector. In practice, I divide the economy into four sectors, so that $S = 4$, which I shall refer to as agriculture, natural resources, manufacturing and services.[8] These four sectors include the entire Canadian economy. As is traditional in the trade literature, the

[8]The agriculture sector consists of the agriculture, forestry, fishing and hunting sectors, which corresponds sector 11 according to the NAICS classification. The natural resource sector consists of mining, quarrying, oil and gas extraction, and utilities, which span NAICS sectors 21 and 22. The manufacturing sector spans NAICS sectors 31-33. Finally, in the service sector I include all service sectors, which include all the sectors above NAICS sector 42 as well as the construction sector, which is sector 23.

service sector, sector 4, is nontraded.

2.4.1 Identification of Prices

In order to operationalize the aforementioned methodology of identifying the impact of global changes on the Canadian economy, I need to separately identify prices and productivities. The production side of the model provides no guidance here, as from the workers' perspective it is only the product of the two that matters. The source of the identification will therefore be the demand side of the model. The first step is to determine the elasticity of substitution between sectors in the CES aggregator, σ, which is taken to equal 2. The next step is to determine the weights of the four sectors in the CES aggregator, $\{\alpha_s\}_{s=1}^4$. In order to do so, I use expenditure shares in Canada in 1992, the first year of my sample and the last year before the start of the simulation. The CES aggregator implies that the expenditure share on sector-s goods is given by[9]

$$shares_{ts} = \alpha_s \left(\frac{p_{ts}}{P_t} \right)^{1-\sigma}$$

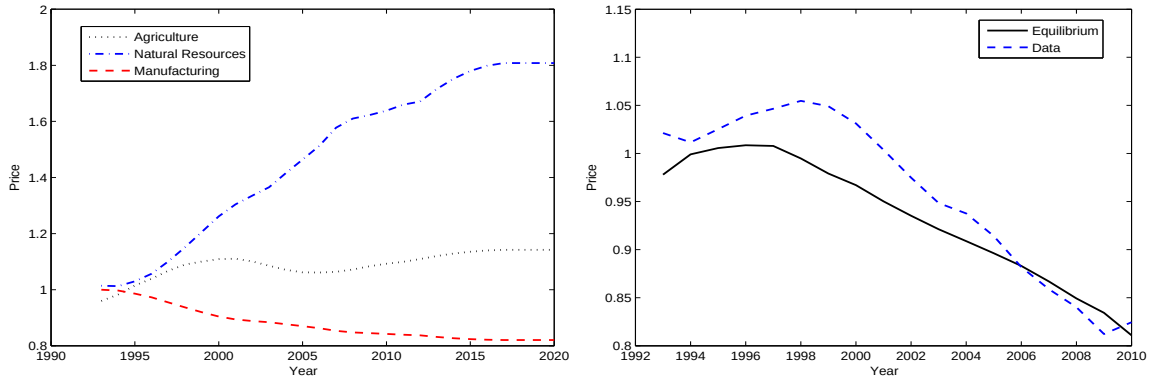
Since I choose units of the goods in each sector so that $p_{1992s} = 1$ for all sectors, the α_s's are equal exactly to the expenditure shares in 1992. This implies α-weights of 0.024, 0.073, 0.163 and 0.740 for agriculture, natural resources, manufacturing and services, respectively.

The above procedure yields all the parameters of the CES aggregator. Given the CES parameters, expenditure shares yield prices in each of the four sectors using the reverse procedure, by virtue of the above equation. Since I use data from 1992 until 2010, this yields prices in each of the traded sectors for those years. I assume that world prices in each of the traded sectors continue to change by the average annual change in the years 1992-2010 in

[9] Expenditure shares are calculated by subtracting net exports from value-added in each sector and dividing by the sum over all sectors of value added less net exports.

each year from 2011 until 2015, and stay fixed at 2015 levels thereafter. The resulting price path is presented in Figure 2.1a.

Figure 2.1: Calibrated Prices



(a) Traded-Sectors (b) Services

Notes: Panel (a) in this figure plots the calibrated traded-sector prices. Traded-sector prices are calibrated directly from the data for the years 1993-2010 and extrapolated from 2011 to 2015. The price of the service-sector good is shown in panel (b). The two plots in that panel represent the service-sector price had that price been calibrated like the other prices, and the service price resulting from the model's equilibrium conditions. Prices are smoothed using a five-year moving average.

Although the above procedure can, in theory, also yield a path of prices for services, I take an alternative route for determining those prices.[10] The reason I do so is that service prices are, unlike the traded-sector prices which are parameters of the model, endogenously determined within the model. As a consequence, exogenously determined service prices would not necessarily constitute equilibrium prices. In order to ensure that the model is in equilibrium, the service-sector prices must be determined endogenously as a solution to the equilibrium conditions of the model. Figure 2.1b plots the path of service prices obtained by the same procedure as used for the traded sectors against the path of service prices that satisfies the model's equilibrium conditions, which, of course, is only derived after

[10] I do, however, use service-sector prices derived in this fashion for parts of the calibration. This is because I must calibrate the model before I derive the equilibrium path of prices in the simulations.

calibration of the entire model and simulation of the entire transition path.[11]

2.5 Calibration

Before moving on to the simulation of the model, the rest of the model's parameters must be calibrated.I begin by specifying the length of a period in the model as corresponding to one week. The choice of a week as the period length compared to, for example, a quarter increases the computation time by a factor 13. However, the length of a period in the model is also the minimum amount of time that a worker must remain jobless after entering a sector, rejecting a job opportunity or dropping out of the labor force. A longer period would make welfare analysis unreliable because it would substantially overstate the welfare loss resulting from unemployment. For similar reasons, an exogenously-imposed long jobless period can significantly distort the decisions of workers as it makes joblessness particularly unattractive. Since the goal of this chapter is to identify how workers' decisions are endogenously affected by changes in the global economy, minimizing the exogenous impact of an arbitrary period length is critical.

The weekly discount factor β is chosen to correspond to an annual real interest rate of 4%, which implies that $\beta = 0.9992$. The exogenous separation rate, λ, is chosen to match the monthly separation rate reported in Shimer (2005) of 0.034 which roughly corresponds to a

[11] Although the trajectory of prices implied by the equilibrium conditions follows the same trajectory as the service prices implied by the data, the equilibrium-implied prices appear to slightly understate the data-implied prices. One potential reason for this is that the service-sector prices are, at least to some extent, determined by world prices. By this interpretation, the difference in the two paths of prices results from an aggregation bias. This is because the service sector might be an aggregate of traded services, whose price is determined by the global economy, as well as nontraded services, whose price is determined by the equilibrium conditions in the local economy. An alternative reason may be that other factors of production may contirbute to the demand for and supply of services. If those factors contribute more to demand for than supply of services, for example, if they are disproportionately represented in production in the traded sectors, then ignoring these factors will lead to an understatement of the equilibrium price of services.

weekly separation rate of 0.008. The exogenous separation rate is assumed to be identical in all four sectors.[12]

Workers live for 50 years, which I map into the data to correspond to ages 20 through 69. Since each period corresponds to one week, this implies that a worker's lifetime, A, spans 2600 periods. The reason that I follow workers into a relatively old age compared with other studies is in order to capture their labor-force-participation decisions. One of the key ways that trade will affect workers is by altering their retirement age, or in the model's language, by altering the labor-market-participation age cutoff. Thus, following workers past the regular retirement age allows me to observe how it changes endogenously, rather than being exogenously determined by the lifespan of the worker.

2.5.1 Human Capital

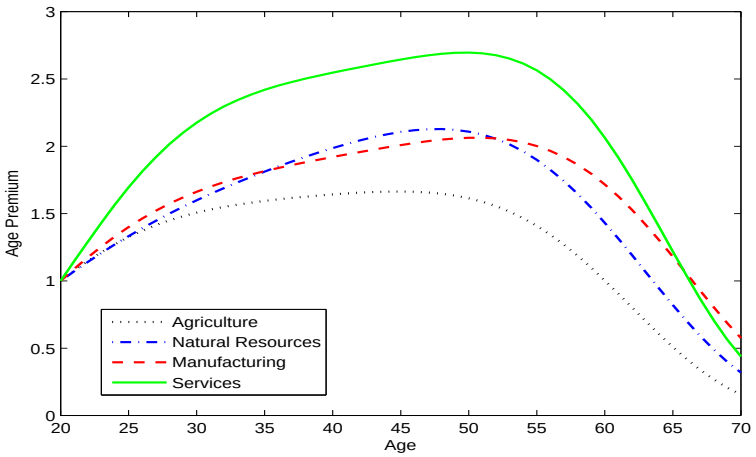
In this subsection, I calibrate the human capital function, $h_s(a,e) = k_s(e)\, g_s(a)$, that relates experience and age to earnings in the four sectors. For this calibration I use data on income by sector and age published by Statistics Canada to obtain an age-earnings profile for each year for each sector. The calibration matches the age-earning profile for each sector averaged over my sample years, 1992 through 2010, which is presented in Figure 2.2.

Because of the strong correlation between age and experience, the empirical literature has not adequately distinguished the effect of age from the effect of experience in the determination of wages.[13] Nevertheless, identifying the separate impact of age is important for understanding how, for example, the retirement age behaves in response to price changes. It is plausible to assume that the upward sloping part of the empirical age-earnings profile is

[12] Attributing the entire separation rate reported by Shimer (2005) to exogenous separations is likely to lead to an understatement of the impact of world prices on the unemployment rate. Since surely some of these separations are, in fact, endogenous, attributing all of them to exogenous separations will lead to an understatement of the number of endogenous separations.

[13] The literature has generally assumed a quadratic, or more recently a quartic, function relating a pseudo-experience variable, defined as age less years of schooling, to the log of earnings. See Murphy and Welch (1990) for a discussion of these issues.

Figure 2.2: Empirical Age-Earnings Profile



Notes: The figure plots the age-earnings profile for workers in each of the four sectors averaged over the sample years 1992 through 2010.

chiefly a result of increasing experience, while the decreasing portion towards the later part of a worker's life is chiefly a result of age-related decreased productivity. Thus, I use the initial portion of the age-earnings profile to identify the impact of experience on earnings and the final portion to identify the impact of age on earnings, as described below.

I assume that $k_s(e)$ takes the following form[14]

$$k_s(e) = k_s(e-1)^{\nu_s} K_s^{1-\nu_s},$$

where $\nu_s < 1$ and $k_s(0) = 1$. This functional form implies a strictly increasing and concave experience-earnings profile that asymptotes to K_s. A higher curvature parameter ν implies a slower convergence to K_s. I calibrate K_s and ν_s to match the earnings growth of workers in the first two decades in which they work, i.e. the initial portion of Figure 2.2.

I assume that the function $g_s(a)$ takes the following form

$$g_s(a) = \begin{cases} 1 & \text{if the worker is less than 50 years old} \\ g_s(a-1) - b_s & \text{if the worker is more than 50 years old} \end{cases}.$$

[14]This functional form is borrowed from Cosar (2010).

Thus, I assume that age begins to affect a worker's productivity at age 50. From age 50 and on their productivity decreases linearly by b_s in every period. I calibrate b_s by matching the earnings growth of workers above the age of 50, given the experience function $k_s(\cdot)$. Thus, taken together the function $k_s(\cdot)$ and $g_s(\cdot)$ are consistent with the empirical age-earning profile when workers are assumed to have experience equal to their age minus twenty.

Since workers produce a quantity equal to $g_s(a) k_s(e) A_s$ and $k_s(0) = g_s(1) = 1$, the productivity, A_s, is properly interpreted as the quantity produced in sector s by a worker up to age 50 with no experience. But, this also means that workers in a sector with a higher A_s are not necessarily more productive than workers in a sector with a lower A_s. Take, for example, the service and agriculture sectors. The return to experience is far greater in the service sector, as is apparent in Figure 2.4a. Thus, given the same A_s in the two sectors, workers in the service sector produce, on average, more than workers in the agriculture sector. As a result, the productivity factors, A_s, which I calibrate next, are not meaningfully comparable across sectors, rather they are informative about changes in productivity occurring over time within a sector.

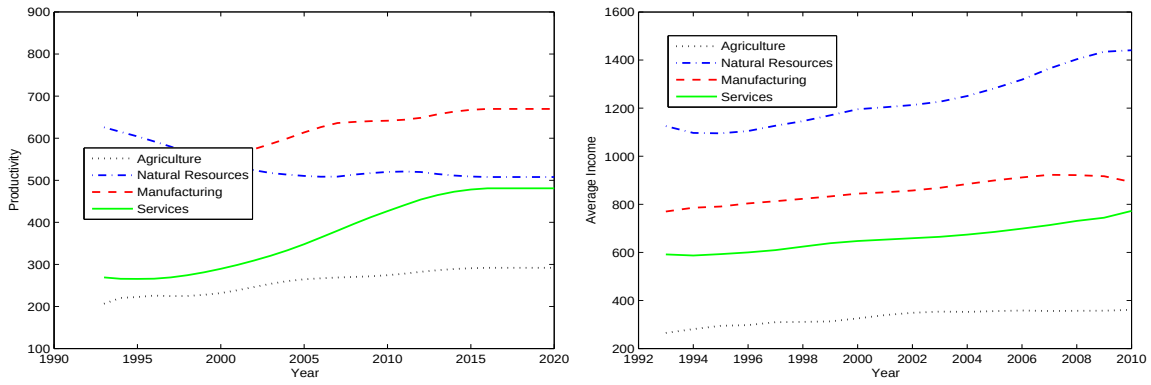
2.5.2 Sectoral Productivities

In this subsection, I calibrate the path of productivities $\left\{\{A_{ts}\}_{s=1}^{4}\right\}_{t=1992}^{\infty}$. I do so by finding for each period t and sector s the productivity A_{ts} such that the average income implied by the the human capital function just calibrated and the empirical distribution of employed workers equals the corresponding average income in the data. Specifically, total output in sector s in period t is

$$Average_Income_{ts} = \frac{p_{ts} A_{ts} \sum_{a=1}^{A} \sum_{e=1}^{A} h(a,e) H(a,e,s)}{\sum_{a=1}^{A} \sum_{e=1}^{A} H(a,e,s)},$$

where $H(a, e, s)$ is the number of employed individuals of age a with experience e in sector s in the data.[15] The numerator is equal to the value added in sector s in period t and the denominator is the number of employed workers in sector s in period t. Using data on the worker distribution and average incomes published by Statistics Canada as well as substituting the price derived in Section 2.4 as p_{ts}, leaves A_{ts} as the only unknown. The solution to the above equation for productivities for each period in each sector is presented in Figure 2.3a.[16]

Figure 2.3: Productivity and Average Income



(a) Productivity

(b) Average Income

Notes: Panel (a) in this figure plots the calibrated productivities. Productivities are calibrated to match average income in each sector for the years 1993-2010, shown in panel (b), and extrapolated from 2011 to 2025. The results are smooted using a five-year moving average.

Before moving on, it is worth considering some of the features of the calibrated productivities. Productivity increases in all sectors except the natural resource sector. A plausible explanation for the decrease in productivity in the natural resource sector is that there is a fixed factor in that sector leading to decreasing returns. For example, as oil production expands, oil companies may search for oil in decreasingly productive locations. However, because of the price increase this remains profitable. An analogous but reverse explanation may explain the increase in productivity in the agriculture sector, which shrunk from 1992

[15] See Appendix B.1 for details on the translation of the empirical distribution to $H(a, e, s)$.

[16] I assume that productivities continue to change at the same average annual rate from 2011 until 2015 as they did from 1992 to 2010. After 2015 I assume that they remain constant until a steady state is reached.

to 2010. Fixed factors are less likely to be as important in the service and manufacturing sectors. It is likely that these sectors experienced a secular increase in productivity unrelated to the increase or decrease in the size of the sector. It is, however, possible that the increase in productivity is a result of access to cheaper intermediate goods caused by the price decline in manufactures, although this is by no means obvious as the opposite may have occurred due to the price increase in the other sectors. Modeling these potential linkages between sector size or intermediate good prices and productivity may lead to important insights that the present model does not capture.

Figure 2.3a and Figure 2.3b differ because average income within a sector depends not only on productivity, but also on compositional effects, workers' labor-market decisions and sectoral prices. Figure 2.3 highlights the importance of disentangling these effects, which the model facilitates, in order to identify the impact of global changes on a small open economy. From Figure 2.3b it appears that manufacturing incomes have increased. But as Figure 2.3a and 2.1a make clear, this increase is, at least in part, because manufacturing productivity increased and in spite of the fact that manufacturing prices decreased. In other words, had manufacturing prices not decreased, manufacturing workers incomes would be even higher. Notwithstanding the complications to the story caused by the compositional effects and workers' labor-market decisions, the underlying intuition for many of the qualitative conclusions from the quantitative analysis are already apparent from the comparison between Figure 2.3 and Figure 2.1a.

2.5.3 Job-Opportunity Distribution

In this subsection, I calibrate the job-opportunity distribution, which consists of the standard deviations of the distribution of job opportunities, η_s, and the disutilities from search and work, c_s. The standard deviation of job opportunities and the disutility from search and work are chosen so that the model's steady-state predictions match the average unemployment

rate and the average labor-force retention rate for workers 55 years old and above, where the averages are taken between the years 1961 and 2010 separately for each of the four sectors. These moments are chosen because the two parameters directly govern the value of searching over dropping out of the labor force and the likelihood that a worker receives an acceptable job opportunity. These, in turn, directly impact the unemployment rate and the labor-force-participation rate, two of the key outcomes of interest.

Given these two parameters, a worker's decisions in sector s can be derived, regardless of the parameters in the other sectors, for any given path $\{z_{ts}\}_{t=t_0}^{\infty} = \{p_{ts}A_{ts}/P_t\}_{t=t_0}^{\infty}$. As a consequence, sector-level unemployment and labor market participation rates in sector s can be computed independently of the other sectors, given a path for z_{ts}. Since prices and productivities have already been calibrated, the path for z_s is already known. In practice, I calibrate the job-opportunity distribution by simulating a steady state with each of the z_s's from 1992 to 2010 and finding the η_s and c_s that ensure that the unemployment rate and labor-force-retention rate for older workers is equal to their corresponding empirical long-run averages. This generates a series of parameters $\{\eta_{ts}\}_{t=1992}^{2010}$ and $\{c_{ts}\}_{t=1992}^{2010}$, the means of which are the parameters used in the calibration.

2.5.4 Sector-Entry Cost Distribution

The final parameters that require calibration are the means and standard deviations, μ_s and θ_s, respectively, of the sector-entry cost draws. I calibrate these by matching the shares of entrants into each of the four sectors predicted by the model in 1992 and 2010, corresponding to the first and last years of my sample, to the actual shares of workers between the ages of 20 and 34 in each of the four sectors in those years. Using two different points in time provides information about the differences in level of the shares of workers entering each sector as well as information about how these shares vary in response to changes in the value of entering a sector. This allows me to pin down both the mean and the standard deviation

of the sector-entry cost draws. These parameters will be critical in determining how changes in world prices affect the relative sizes of the sectors.

Since, at this stage, all the parameters have been calibrated including prices and productivities from 1992 and on, it is straightforward to compute the value for a worker of entering sector s in any year t, V_{ts}.[17] Given these values of entering each sector, I can compute the share of workers that enter each sector in each year for any set of parameters $\{\mu_s, \theta_s\}_{s=1}^4$. The calibration then consists of finding a set of parameters such that the share of entrants predicted by the model in 1992 and 2010 matches the empirical share of workers between the ages of 20 and 34, who serve as a proxy for entrants, in those two years.

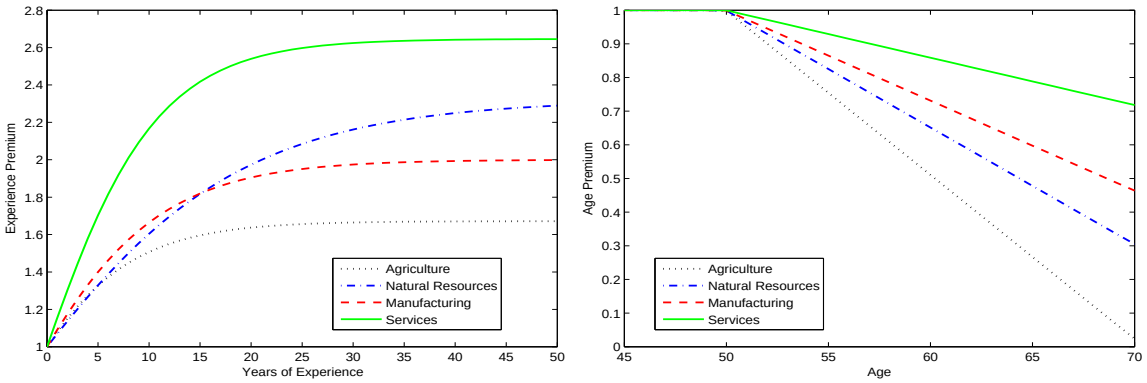
2.5.5 Calibration Results and Discussion

The results of the calibration of the functions governing productivity by experience and age, $k_s(\cdot)$ and $g_s(\cdot)$, are presented in the two panels of Figure 2.4. Experience is most important in the service sector and least important in the agriculture sector. A worker with twenty years experience in the service sector is approximately two-and-a-half times more productive than a worker with no experience, whereas in the agriculture sector such a worker is only slightly more than one-and-a-half times more productive than a worker with no experience. This is important because one of the incentives for accepting a job is as an investment in human capital or experience. In fact, I find that, in the service sector, some workers take jobs that offer negative periodic utility because work experience is so valuable in that sector.

It is reasonable to expect that in sectors that require more physical labor and heavy lifting the age effect on productivity would be particularly strong. This expectation appears to be borne out in the calibration results. Services is likely the sector that requires the least physical exertion, while agriculture likely requires the most. Indeed, as Figure 2.4b shows,

[17] See Appendix B.2.1 for a description of this computation.

Figure 2.4: Productivity by Experience and Age



(a) Experience (b) Age

Notes: Panel (a) of the figure charts the relationship between a worker's experience and her productivity. Panel (b) charts the decrease in worker productivity once she passes the age of 50.

service sector workers experience only a small drop off in productivity as they age, while agriculture workers are almost entirely unproductive by the time they reach their seventies.

The results of the calibration of the job-opportunity and sector-entry-cost draws are summarized in Table 2.1. Since I match the sector-entry cost distributions to the shares of entrants into each sector, I can only identify three distributions. The fourth share of entrants is uniquely determined by the other three. The sector-entry-cost draws should therefore be interpreted as the cost of entering a sector relative to the cost of entering the service sector. The mean cost of entering the non-service sectors relative to the service sector is positive for all sectors. This is not surprising given the fact that the service sector attracts more than 80 percent of entrants.

The agriculture sector carries with it the lowest disutility from work-related activities, c_s. The reason for this is that incomes are so low in the agriculture sector that if the work carried with it more disutility, then nobody would ever even bother to search for work in that sector. Both the disutility cost of work, c_s, and the cost of entry, μ_s, play an important role in explaining the pattern of sectoral entry vis-a-vis lifetime earnings in a sector. Lifetime earnings are highest in the natural resource sector followed by the manufacturing sector.

Table 2.1: Second-Stage Calibration

	Agriculture	Natural Resources	Manufacturing	Services
c_s	0.19	0.93	0.85	0.73
η_s	0.10	0.17	0.20	0.12
μ_s	0.79	2.46	0.35	0
θ_s	0.32	1.00	0.20	0

This table presents results from the second-stage calibration. c_s is the utility cost of work-related activities. η_s is the standard deviation of job-specific disutility draw. μ_s is the mean of the idiosyncratic sector-entry cost draw. θ_s is the standard deviation of the idiosyncratic sector-entry cost draw. In all cases, the subscript s denotes the sector.

Yet, most entrants choose to work in the service sector. This is because, as can be seen in Table 2.1, the natural resources and manufacturing sectors are far more costly both to enter and to work in relative to the service sector.

Finally, the standard deviations of the sector-entry-cost draws, θ_s, are key in determining how much a sector grows when it experiences a positive shock, or vice versa. Indeed, these standard deviations are identified through the change in sectoral entry patterns vis-a-vis the change in the value of entering each sector between 1992 and 2010. As a result, θ_s in the natural resource sector is relatively high, whereas it is relatively low in manufacturing. The reason is that the natural resource sector grew only by a little even though the value of entering that sector increased tremendously from 1992 to 2010, while the manufacturing sector shrunk by much more even though the decrease in entry value was much smaller compared with the natural resource sector.

2.5.6 Initial Worker Distribution

The final issue that I address in this section is the determination of the initial worker distribution. In order to simulate the model, an initial worker distribution is necessary. This

distribution is important in order to account both for the compositional effects and the general equilibrium effects in the model. Since my sample begins in 1992, I initialize the simulation, which begins in 1993, with a worker distribution matching the empirical distribution in 1992 based on labor-force statistics published by Statistics Canada. This data provides information on the size of the labor force and the number of employed and unemployed workers by age and sector. I assume that a worker's years of experience is equal to her age minus 20.

The main issue that arises in the calibration is that, for employed workers, there is no direct observable counterpart in the data to the job-specific disutility, j, postulated by the theoretical model. I assign jobs to employed workers as follows. Using the parameters found in the calibration, I compute the job-offer-acceptance cutoffs for each age and experience level in 1992. I then assign employed workers at each age and experience level to jobs according to the probability of getting that job opportunity conditional on the job opportunity being above this cutoff. The details of the remainder of the determination of the initial worker distribution is relatively mechanical and is therefore relegated to Appendix B.1.

2.6 Quantitative Analysis

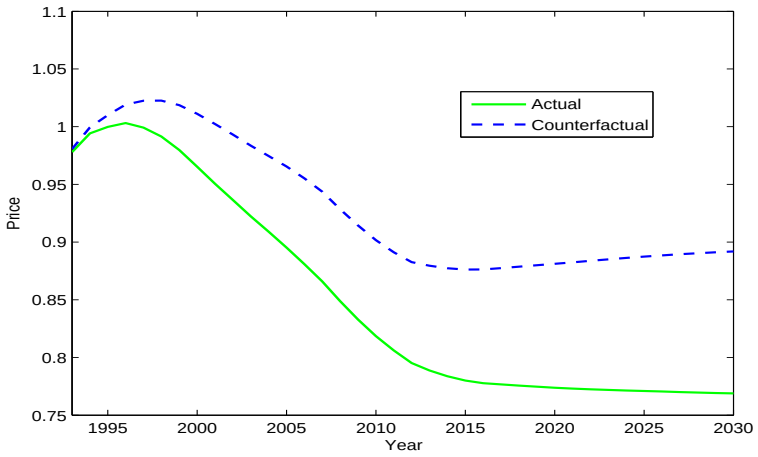
My goal in this section is to provide quantitative estimates of the impact of changes in world prices on the Canadian economy and particularly on Canadian workers. To reiterate the logic underlying the analysis, I first solve the model with the actual traded-sector prices identified in Section 2.4. I then solve the model with traded-sector prices fixed counterfactually at 1992 levels. The differences in outcomes between the two simulations are interpreted as the impact of changes in the global economy on the Canadian economy from 1993 onward.[18]

Since I do not model the global general equilibrium that determines world prices in the traded sectors, I cannot attribute my findings to any specific cause. This is the reason that I will refer to the impact of global changes or alternatively changes in world prices, rather

[18] In order to solve the model, I employ a shooting algorithm, the details of which are in Appendix B.2.

than something more specific such as the growth of China or other developing countries. Nonetheless, it is worth noting that the price changes observed are, at least qualitatively consistent with how we might plausibly expect productivity growth in the developing world to affect world prices. On the one hand, the decrease in the price of manufactures is consistent with increased productivity and import penetration from China and other developing countries. On the other hand, the increase in the price of natural resources is consistent with the increased demand for energy by these countries. Thus, although the model does not identify them as such, it is reasonable to interpret the results as, at least in part, the effect of productivity growth in the developing world on Canadian workers.

Figure 2.5: Equilibrium Service Prices



Notes: The figure plots the path of service sector prices that constitute an equilibrium in the actual and counterfactual simulations. The plots are smoothed using a five-year moving average.

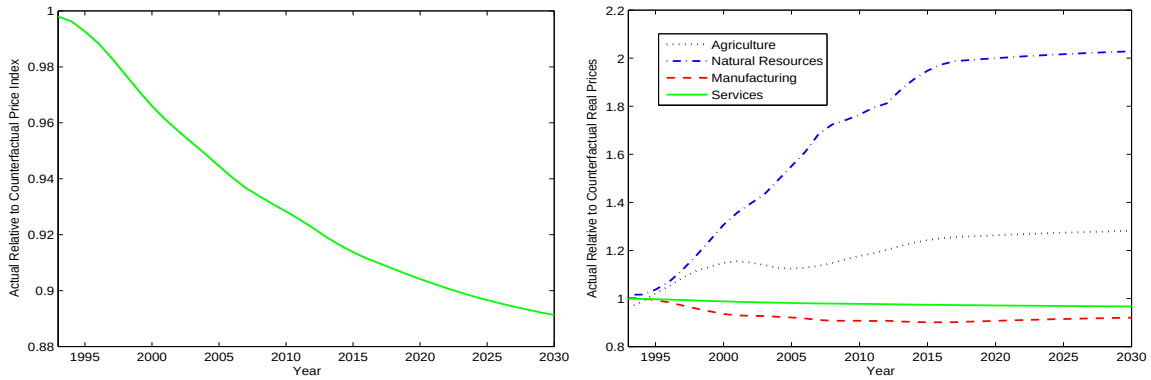
The equilibrium path of service prices, which constitute a solution to the equilibrium conditions outlined in Section 2.2.3, are presented in Figure 2.5. Both paths exhibit a general downward trajectory. The reason is that in both cases the service sector is growing which increases the supply of services, thereby exerting downward pressure on the price. However, in the case of the counterfactual simulation the price does not drop by nearly as much as in the actual simulation. This is because global changes lead to further growth in the service sector, which would not have occurred in the absence of these changes. The mechanisms that lead to this growth in the service sector will become clearer as we delve

deeper into the results of this section.

2.6.1 Prices

Workers' decisions depend directly on the price of the good they produce relative to the price of the good they consume. As a result much of the following analysis will hinge directly upon how changes in relative world prices affect the price of goods in each sector relative to the price of the final good. I therefore begin the analysis with a description of how global changes affect both the price index and good prices relative to the price index, or the real prices. These are presented in Figure 2.6.

Figure 2.6: Price Index and Relative Prices



(a) Price Index (b) Real Prices

Notes: Panel (a) of the figure charts the impact of global changes on the price index from 1993 to 2030. Panel (b) charts the price of the sectoral goods relative to the price of the final good in each of the four sectors from 1993 to 2030. Both panels are smoothed using five-year moving averages.

Global changes lead to a decrease in the price index relative to the counterfactual with no global changes. In fact, by 2030 global changes account for more than a ten percent decrease in the cost of the final good. Although the world prices of agriculture and natural resources increase by more than either prices in the manufacturing or service sectors, as can be seen from Figure 2.1a and 2.5, the net effect is a decrease in the price index. The reason

is that services and manufacturing carry far more weight in the CES aggregator than do the other sectors, and the decrease in their respective prices outweighs the larger increases in the prices in the other sectors.

Figure 2.6b shows that global changes account for almost a doubling of the real price of natural resources by 2030 and to a smaller but still substantial increase for the agriculture sector. The situation is reversed for the manufacturing and service sectors. A 10 percent decrease in the real price of manufactures is attributable to global changes, while there is a smaller but still negative effect for the service sector.

The reason that there is such a small effect on the service sector is that it carries almost three quarters of the weight in the CES aggregator. As a result, the price index tracks the movement of service prices very closely, shielding service-sector workers from large changes in the real price of the good they produce. Indeed, this is the reason that service-sector workers are least affected by fluctuations in world prices, as we shall see below.

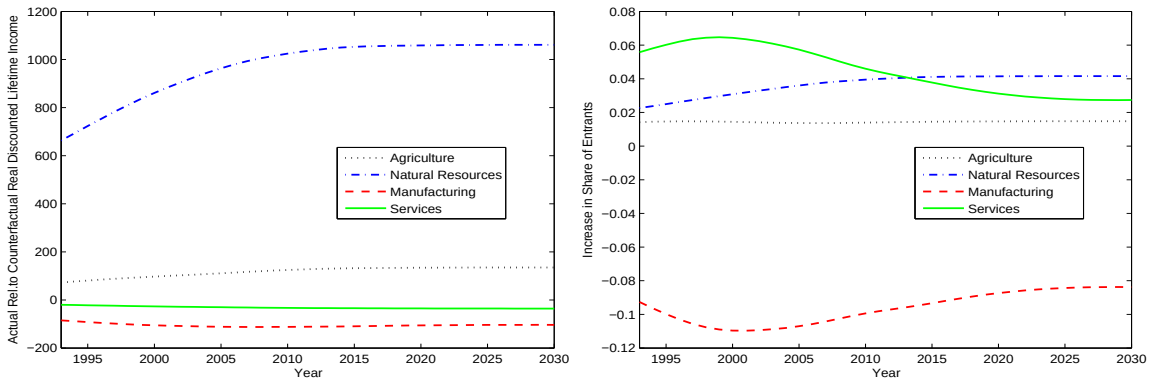
2.6.2 Sectoral Choice and General Equilibrium Effects

When relative world prices change, so too does the relative attractiveness of the different sectors. This leads to adjustments in workers' sectoral choices, one of the key mechanisms through which the economy adjust to changes in relative prices. However, adjustment in sectoral choice then feeds back into relative prices through changes in both the demand for and supply of the service-sector good, as discussed in detail in Section 2.3.3. Therefore, a fuller understanding of the price dynamics discussed in the previous subsection requires an analysis of the impact of global changes on the dynamics of sectoral choice in the economy.

The determinant of the shares of young workers that enter each sector is the expected discounted lifetime utilities afforded to the worker in each of the four sectors. Figure 2.7a presents the increase in expected discounted lifetime real income attributable to changes in

world prices for workers entering the labor market from 1993 to 2030.[19] The effect of global changes on real lifetime incomes is in line with its effects on relative prices considered in the previous subsection. However, the effect on shares of entrants does not follow the same pattern. Figure 2.7 shows that, although global changes lead to a decrease in real lifetime income in the service sector, they nevertheless lead to an increase in the share of entrants. World price changes lead to a 25,000 decrease in lifetime income in the service sector for workers born in the year 2000, and yet these changes also lead to a 6 percentage point increase in the number of entrants. In the natural resource sector by 2020 global changes are responsible for an increase in the real lifetime income of new entrants of approximately 1,000,000 dollars, and yet these same global changes are only responsible for a 4 percentage point increase in entrants into that sector.

Figure 2.7: Discounted Lifetime Earnings and Sectoral Entry



(a) Lifetime Income

(b) Entrants

Notes: Panel (a) plots the absolute increase in expected discounted lifetime income due to changes in world prices in each of the four sectors for newborn workers from 1993 to 2030. The changes are in thousands of 1992 Canadian dollars. Panel (b) plots the change in shares of entrants into the four sectors from 1993 to 2030 attributable to changes in the global economy. The plots are smoothed using a five-year moving average.

Why do more young workers enter the service sector despite the decrease in lifetime earnings? Why do more workers not enter the natural resource sector despite the massive

[19] In this case, I present the absolute increase rather than the relative increase because it is the absolute change in income that matters to the worker for her sectoral choice. In addition, I present the change in income rather than changes in utility as these are almost perfectly correlated, and the former is more easily comparable to observable data.

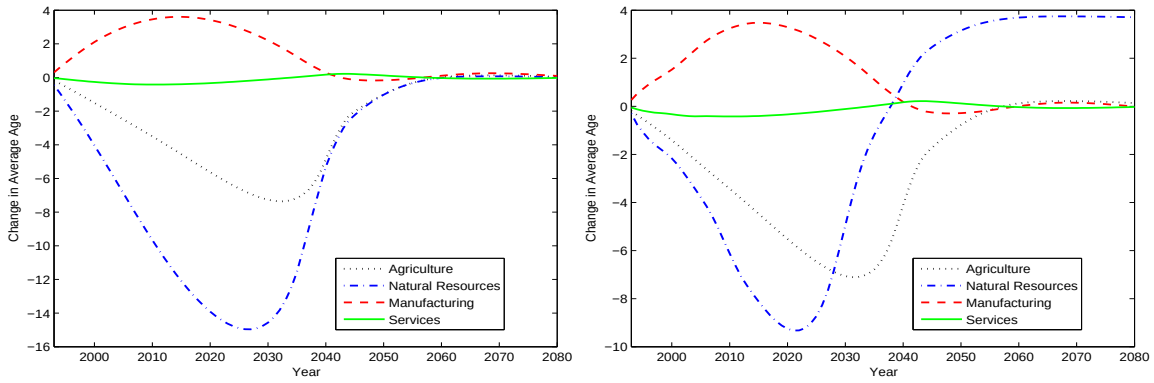
increase in their lifetime earnings? The answer to these question lies in the curvature of the distribution of the sector-entry cost draws, determined by their standard deviations. The standard deviation of the cost draws in the agriculture sector and especially in the natural resource sector is far higher than in the other sectors, as reported in Table 2.1. As a result, those sectors require a much higher increase in lifetime income to attract more workers compared with the service sector. Therefore, a worker who would have entered the manufacturing sector in the absence of changes in world prices is now more likely to enter the service sector rather than the other sectors. The service sector does, after all, become more attractive *relative* to the manufacturing sector. Moreover, the service sector requires a smaller increase in lifetime income relative to the manufacturing sector to attract workers compared to the increase that would be required of the natural resource or agriculture sectors.

From the results presented thus far, the following picture emerges. Due to changes in the global economy the real price of manufactures decreases. As a result, more workers shun the manufacturing sector and instead choose other sectors, especially the service sector. The growth in the service sector and the resulting increase in the supply of services exerts downward pressure on the price of services. Eventually, however, the gap in attractiveness between manufacturing and services narrows due to the downward pressure on service prices. This can be seen in Figure 2.7b with the reversal in the entry pattern beginning around the turn of the millennium. This is the general-equilibrium effect in action.

2.6.3 Compositional Effects

Before moving on to a discussion of the impact of global changes on labor-force outcomes, it is useful to consider how these changes affect the average age within each sector. Indeed, these compositional effects will prove indispensable in properly interpreting the impact of global changes on Canadian labor-market statistics. Figure 2.8 plots the increase in average age within each of the four sectors caused by changes in world prices.

Figure 2.8: Increase in Average Age Attributable to Changes in World Prices



(a) Population (b) Labor Force

Notes: The figures chart the change in average age in the population and labor force in each of the four sectors from 1993 to 2080. The plots are smoothed using a five-year moving average.

In the long run the change in world prices has no effect on the average age of the population within each sector. The reason is that in the steady state the same number of entrants enter in every period regardless of the size of the sector, so that there are no compositional effects. However, in the short run, sectors that shrink relative to the counterfactual experience an increase in their average age compared to the counterfactual, as is apparent in the manufacturing sector. The opposite is true for the other sectors. The decrease in average age in the service sector is barely perceptible because the size increase attributable to global changes is small relative to the original size of the sector.

The average age of the labor force displays somewhat different dynamics. The reason is that the average age is affected both by the average age in the population and the average age at which workers drop out of the labor force. As a result, unlike in the case of the population average age, even in the long run the change in world prices has an effect on the average age in the labor force. This is most clearly visible for the natural resource sector. The change in world prices induces workers to stay in the labor force for longer. This effect will tend to increase the average age in the labor force. However, the price change also induces new entrants to enter the sector. This effect will tend to decrease the average age of the labor force. The first effect is permanent, but the second is only temporary. As a result, initially

the second effect dominates, but as it tapers off the first effect takes over and leads to an increase in the average age of the labor force even in the long run.

Why is the first effect almost nonexistent in the other sectors? In the agriculture sector, the reason is that the drop off in productivity as workers age is so great that even the increased real price of agricultural goods does not induce a very large increase in labor force participation among older workers. For the services and manufacturing sectors the reason is different. Both those sectors experience an increase in their income in both the actual equilibrium and the counterfactual equilibrium, albeit a larger increase in the counterfactual. Nevertheless, even the smaller increase in the actual equilibrium suffices to make even the oldest workers participate in the labor force. As a consequence, the price change does not induce a change in labor market participation among older workers.[20]

2.6.4 Labor-Market Outcomes

The purpose of the results so far has been to paint a broad picture of the effect of changes in world prices on the Canadian economy, and to provide the tools to understand how these changes affect labor-market outcomes. We are now ready for the main goal of this section, namely to quantify the impact of changes in the global economy on labor-market outcomes of workers in Canada. In particular, I will focus on the impact of world prices on labor-force participation, unemployment and income.

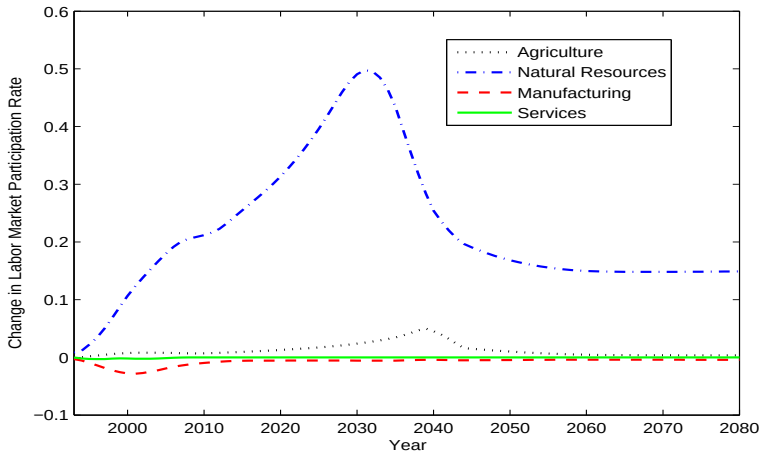
Labor-Force Participation

From Proposition 2.3.1 we can conclude that in sectors in which global changes led to higher prices, these same changes also caused a decrease in the labor-force-participation age cutoff. Proposition 2.3.2 guarantees a higher-labor-force-participation rate in the steady state of the

[20] If workers were allowed to work to an arbitrarily old age, then the effect would return. It is the fact that they are in a corner solution at age 70 that negates an increase in labor-market participation in the counterfactual equilibrium.

actual equilibrium compared to the counterfactual equilibrium for sectors that experience a price increase. However, along the transition path, this is not guaranteed because it depends on the initial worker distribution. Nevertheless, as Figure 2.9 reports, it turns out that for the Canadian economy after 1992, sectors that experienced a price increase, in fact also experienced an increase in labor-market-participation rates attributable to changes in world prices, and vice versa.

Figure 2.9: Change in Labor-Force-Participation Rate Attributatble to Changes in World Prices



Notes: This figures shows the changes in labor-market-participation rates from 1993 to 2080 attributable to changes in world prices. The plot is smoothed using a five-year moving average.
.

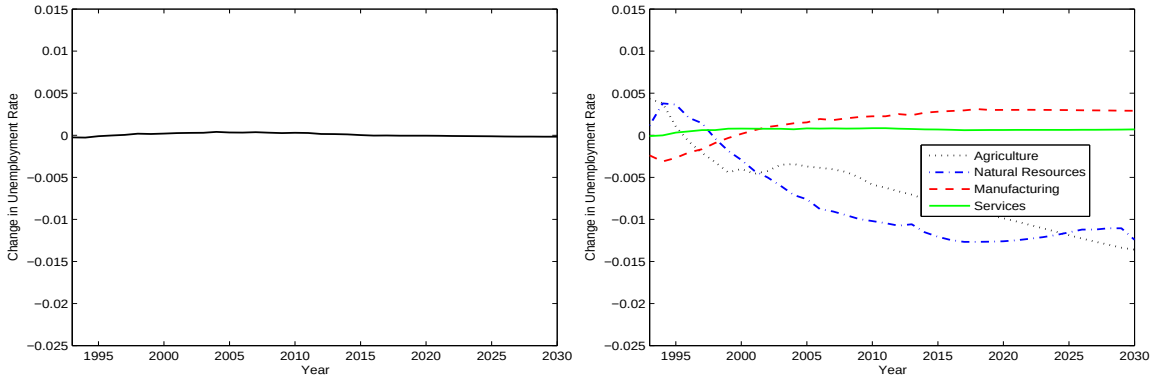
Although labor-force participation does follow the expected qualitative pattern, it does not do so in a monotonic fashion. The reason is compositional effects. In fact, the change in the labor-force-participation rate attributable to changes in world prices follows the opposite pattern as the change in average age attributable to changes in world prices. This is because it is the older workers who drop out of the labor force. As a result, a decrease in average age results in an increase in labor-force participation. The reason that the change in labor-force-participation rate attributable to the world price change reverts to zero for the manufacturing sector in the early 2000's is that, as mentioned above, in both the actual and counterfactual equilibria everyone participates in the labor force. Thus, there is no effect in

that sector on labor-force-participation after approximately 2015.

Unemployment

How have global price changes affected unemployment in Canada? Figure 2.10a shows that changes in world prices have had little appreciable effect on aggregate unemployment. But Figure 2.10b shows this conclusion to be a misleading result of aggregating across sectors. Indeed, on the one hand, in the long run, global changes are responsible for a decrease in the unemployment rate of approximately 1.3 percentage points in the natural resource and agriculture sectors. On the other hand, in the long run, global changes are responsible for a 0.3 and 0.1 percentage point increase in the unemployment rate in the manufacturing and service sectors, respectively.

Figure 2.10: Increase in Unemployment Rate Attributable to Changes in World Prices



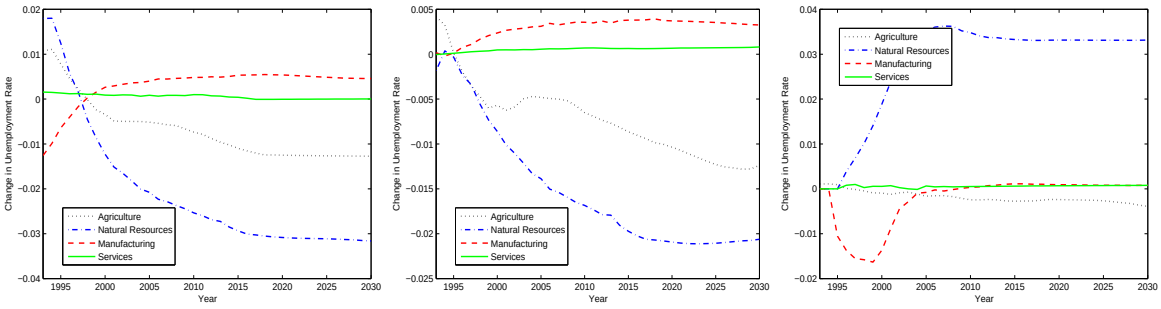
(a) Total (b) By Sector

Notes: The two figures show the change in the unemployment rate from 1993 to 2030 attributable to changes in world prices. Panel (a) is an aggregate over all workers and panel (b) decomposes the aggregate into the four sectors. Both plots are smoothed using a five-year moving average.

However, focusing only on the long-run effect of global changes misses some of the interesting short-run dynamics that can shed further light on how these changes affect the Canadian economy. Indeed, Figure 2.10b shows that in the short run unemployment moves in the opposite direction as compared to the long run in all four sectors. Changes in the

real price in a sector affects the unemployment rate in that sector in three ways. First, as per part one of Proposition 2.3.1, the job-opportunity acceptance rate moves in the same direction as the price change. Thus, every search period is more likely to result in an employment spell when the real price of a good increases, and vice versa for a price decrease. This effect will tend to make the unemployment rate move in the opposite direction of the price change. Second, as per part two of Proposition 2.3.1, the likelihood that an old jobless worker searches rather than retires moves in the opposite direction of the price change. This will tend to make the unemployment rate change in the same direction as prices. Third, since new entrants into a sector are unemployed, the compositional effect will tend to increase unemployment in growing sectors because they have a higher proportion of new entrants relative to incumbents.

Figure 2.11: Change in the Unemployment Rate Attributable to Changes in World Prices by Age



(a) Age 20 to 25 (b) Age 30 to 60 (c) Age 65 to 70

Notes: The three figures show the evolution of the change in the unemployment rate attributable to changes in world prices for different age groups from 1993 to 2030. The plots are smoothed using a five-year moving average.

To hone in on these three effects, Figure 2.11 decomposes the change in the unemployment rate attributable to world price changes for the four sectors into the changes in the unemployment rates among young, middle-aged and old workers. I will describe the effect of a price change for the manufacturing sector, but an analogous, if opposite effect, holds for the other sectors. Since the manufacturing sector is shrinking due to the decline in the real price of manufactures, the compositional effect implies that initially new entrants

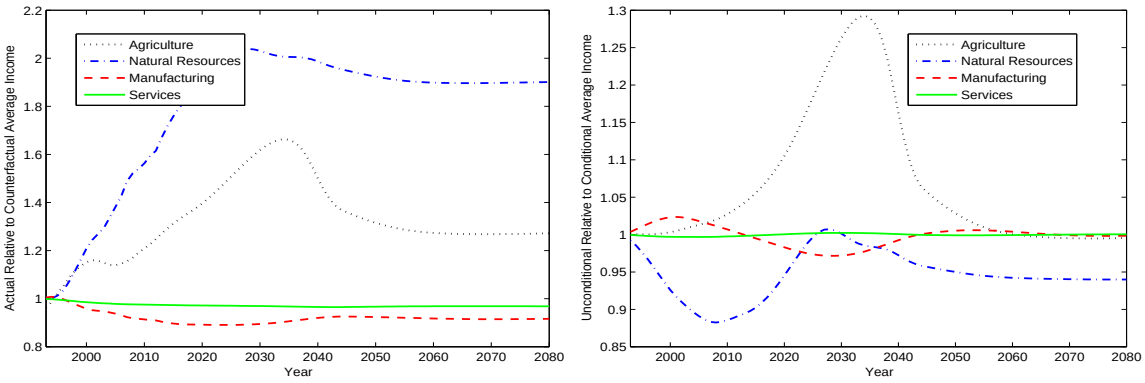
are underrepresented in the population. Since these workers have the highest unemployment rate, this effect leads to lower unemployment in the manufacturing sector. As more similarly sized cohorts enter, this effect is diminished and eventually is dominated by the job-opportunity-acceptance effect. This is the story told by Figure 2.11a.

For middle-aged workers only the job-opportunity-acceptance effect is relevant. Indeed, 2.11b shows that the price decline in manufactures is responsible for an increase in the unemployment rate among middle-aged workers along the entire transition path. Unlike for young and middle-aged workers, for old workers the labor-force-participation effect comes into play. Figure 2.11c shows that for old workers in the manufacturing sector, the price decline initially leads to a decrease in unemployment. This is because the price decline induces some workers to drop out of the labor force rather than search. Eventually, however, this effect disappears, because, as I have already mentioned, after the early 2000's the price decline does not induce manufacturing workers to retire early. When that effect disappears it is again the job-opportunity-acceptance effect that dominates. However, in the natural resource sector this effect remains during the entire transition path and dominates the job-opportunity-acceptance effect among older workers. The sum of all these effects results in the sector–level unemployment effects presented in Figure 2.10b.

Income

The final labor-market outcome that I consider is income. Figure 2.12a reports the effect of changes in world prices on average income within each sector without controlling for age and experience. Indeed, when controlling for age and experience the change in average income caused by global price changes is equal to the change in real prices, presented in Figure 2.6b. This is because a worker's real income is equal to $h(a,e) A_{ts} p_{ts}/P_t$. Since A_{ts} does not change between the actual and counterfactual simulations, relative real income conditional on a and e is equal to the relative real price.

Figure 2.12: Average Income



(a) Unconditional (b) Unconditional Relative to Conditional

Notes: Panel (a) plots the effect of global changes on the unconditional average income in each sector from 1993 to 2080. Panel (b) plots the the effect on average income conditional on age and experience over the same time period. Both plots are smoothed using a five-year moving average.

From a worker's perspective, it is the change in income conditional on her age and experience that matters. The compositional effect, which leads to the divergence of the conditional and unconditional average incomes, is of no consequence to the individual worker. The ratio of the change in the unconditional average attributable to the changes in world prices and the change in the conditional average attributable to changes in world prices is presented in Figure 2.12b. This ratio is, in a sense, a measure of the overstatement of the change in average income that results from failing to account for compositional changes.

To make sense of Figure 2.12b, consider again the manufacturing sector. The real-price decrease in that sector is responsible for an decrease in the flow of entrants compared to the counterfactual with no price increase. In the first stage, the new entrants earn, on average, less than the sector's incumbents, as per the age-earnings profile reported in Figure 2.2. As a result, this leads to the initial increase in Figure 2.12b since the new entrants are underrepresented in the manufacturing sector relative to the incumbents. In the second stage, as the average age of those induced by the price changes to enter the sector increases, so too does their average income. In contrast, as the average age of the incumbents increases, eventually their average income decreases as a larger proportion of the incumbents faces the

downward sloping segment of the age-earnings profile. The increase in average income of the new entrants together with the decrease in the average income of the incumbents induces the reversal apparent in Figure 2.12b. Eventually, the compositional effect actually leads the unconditional average income to understate the conditional average income. This is because now the underrepresented entrants earn more, on average, than the overrepresented incumbents. In the final stage, the compositional effect disappears as the size of the dying cohorts is similar to the size of the entering cohorts.

The effect is reversed in the service and natural resource sectors in which the change in world prices induces increased entry. Like in the case of the other labor-market outcomes the effect on the service sector is small both because the real price change induced by the change in world price is small for this sector and because the increased entry induced by the change in world prices is small relative to the initial size of the sector. In the agriculture sector the upward-sloping portion of the age-earnings profile is relatively flat, whereas the downward sloping portion is much more pronounced than in the other sectors. As a result only the final two stages of the evolution of the composition effect are apparent in that sector.

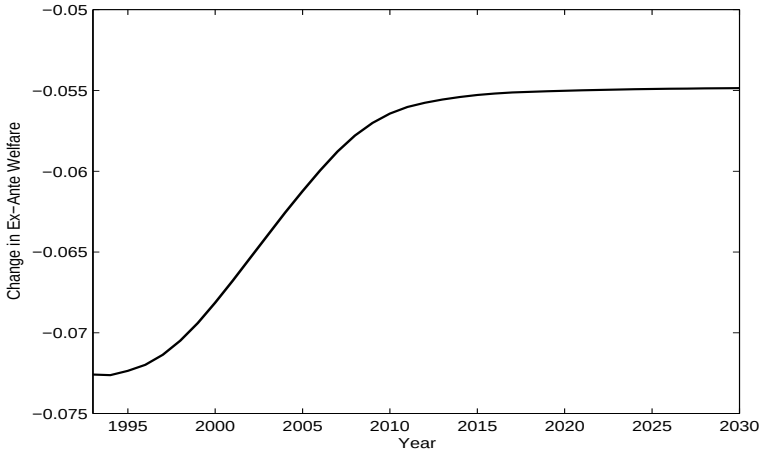
Finally, in the natural resource sector there is an additional effect because the age of the labor force increases due to the increase in the real price of natural resources. Since old workers induced by the price change to work earn less than the average in the sector, this effect leads to a decrease in the unconditional relative to conditional average incomes. This effect, unlike the compositional effects just discussed, is permanent. As a result, even in the long run the change in average income induced by the price change in the natural resource sector is an understatement of the change in the conditional average income.

2.7 Welfare

In this section, I turn to what are in a sense the most important results in this chapter, namely the welfare effects of changes in world prices. How do all the quantitative results presented

in the previous section translate into welfare outcomes for Canadian workers? Figure 2.13 shows that the changes in world prices are responsible for approximately a 5.5% decline in aggregate welfare in the long run. That is, the expected discounted utility of a young worker, prior to drawing sectoral-entry costs, is 5.5% lower in the actual equilibrium as compared to the counterfactual equilibrium. Along the transition path this loss is even bigger, reaching almost 7.5% for workers born in the mid 1990's. The welfare loss is decreasing over time mostly because of the ongoing increase in real prices in the natural resource sector, which mitigates some of the negative impact of the changes in the price of manufactures and services.

Figure 2.13: Changes in Aggregate Welfare Attributable to Changes in World Prices



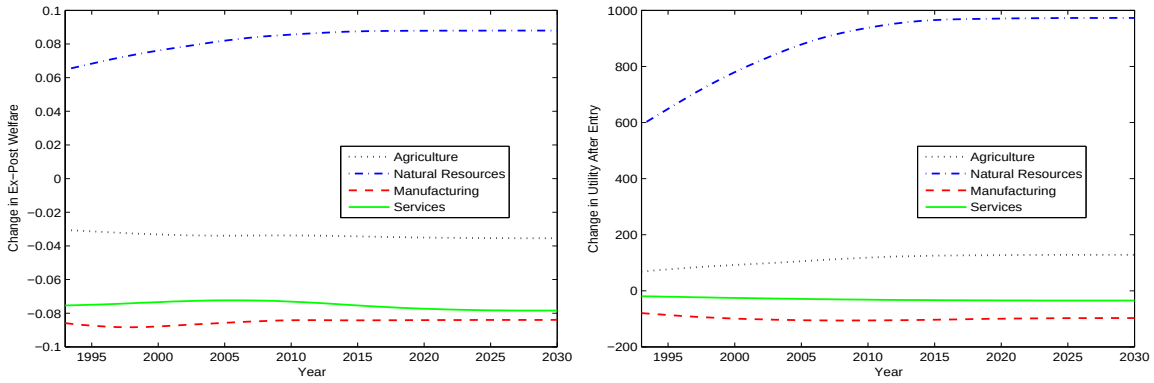
Notes: The figure plots the aggregate change in welfare attributable to changes in world prices. This is the same as the change in expected discounted lifetime utility of a worker prior to drawing sector-entry costs.

To fully understand the decrease in aggregate welfare, it is instructive to decompose the welfare effect of changes in world prices by workers' sectoral choice. Figure 2.14a reports the change, attributable to changes in world prices, in the average lifetime discounted utility of workers in each of the four sectors. This change is the sum of two effects. First, it depends on the change in the average lifetime discounted utility of a worker after entering a sector, presented in Figure 2.14b. This effect will tend to move average lifetime utility of a worker in a sector in the same direction as the change in real price in that sector. Second, it depends

on the average cost of entry incurred by an entrant. This effect will tend to decrease average lifetime utility in growing sectors and increase it in shrinking sectors.

Both these effects conspire to decrease the average lifetime utility in the service sector. However, in the other sectors the two effects operate in opposite directions. Only in the agriculture sector does the second effect dominate, leading to a decrease in the average lifetime utility of agriculture workers despite the increase in post-sectoral-entry utility. Nevertheless, those workers who would have entered the agriculture sector even without the changes in world prices, i.e. those with the lowest entry costs, do indeed benefit from the change in world prices.

Figure 2.14: Changes in Welfare Attributable to Changes in World Prices



(a) Total (b) By Sector

Notes: Panel (a) reports the percentage change in average lifetime utility of a worker in each of the four sectors accounting for the sector-entry cost. Panel (b) shows the increase in expected discounted lifetime utility for a worker who has just entered a sector and already incurred the entry cost.

The standard deviations of the sector-entry-cost draws are the key determinants of how sector entry patterns respond to changes in the expected post-entry utilities offered by each sector. As such, they are also important in determining the general-equilibrium effects discussed in Section 2.3.3. These general-equilibrium effects, in turn, affect the post-entry utilities. Indeed, if all sectors are traded, then these general equilibrium effects would be nonexistent, and, as a result, the parametrization of the sector-entry draws would have no effect on post-entry utilities. Similarly, other sector-level statistics, such as unemployment

and labor-force participation, depend on sectoral entry patterns only through their effect of service prices.

Even in the absence of general-equilibrium effects, sectoral entry patterns affect aggregate statistics because they determine the relative sizes, and, as a result, the relative weights placed on the various sectors in computing aggregate outcomes. Thus, the parametrization of the sector-entry cost draws are important in determining the aggregate welfare effects of changes in world prices. In most models, the change in prices would lead to a mass exodus of workers from the manufacturing and service sectors to the agriculture and natural resource sectors. If this were the case aggregate welfare may well have increased in response to the change in world prices. However, these models fail to account for the fact that young entrants are not so easily swayed to change their sectoral choice. The loss in welfare in Canada comes about precisely because those sectors that gain from the changes in world prices are also those that find it hardest to attract new workers.

2.8 Conclusion

The goal of this chapter has been to study the impact of changing world prices on a small open economy. In order to facilitate the analysis, I build a dynamic general-equilibrium model of a small open economy with overlapping generations. The key innovation in the model is that it simultaneously incorporates endogenous sectoral entry patterns, unemployment and labor-market participation in a dynamic setting. I derive some of the models qualitative properties, for example, that sectors that experience a price increase experience an increase in labor force participation and a decrease in the jobless rate. However, these results mainly underscore the need for a quantitative analysis of the impact of changes in world prices on a small open economy. Indeed, such a quantitative analysis is the main contribution of this chapter.

To facilitate the quantitative analysis, I begin by calibrating the model to match important features of the Canadian economy. Using the calibrated model, I first simulate the economy's performance with the actual time series of world prices from 1993-2010. I then simulate the economy's performance with world prices counterfactually held fixed at their 1992 levels, but with other parameters remaining unchanged from the previous simulation. The resulting differences between the two simulation is interpreted as the impact of changes in world prices on the Canadian economy. Importantly, unlike other quantitative analyses of trade shocks, I do not assume that the Canadian economy was in a steady state prior to the trade shock or that productivities do not change throughout the economy's transition path.

I find changes in world prices are responsible for an initial decline in unemployment in the manufacturing sector and to an initial increase in the agriculture and natural resource sectors. However, this is mostly due to transitional compositional effects. Indeed, in the long run, a 0.3 and 0.1 percentage point increase in the unemployment rates in the manufacturing and service sectors, respectively, are attributable to changes in world prices, while these same changes are responsible for over a 1 percentage point decline in the unemployment rate in the remaining sectors. Nevertheless, because of the relative sizes of the sectors, aggregate unemployment remains almost completely unchanged along the economy's entire transition path. Furthermore, I find that labor-force participation increases substantially in the natural resource sector, but is almost completely unaffected in the other sectors.

I find that some of the structural change in the Canadian economy is attributable to changes in world prices. In particular, approximately a 5 percentage point decline in the share of workers entering the manufacturing sector is attributable to the decline in the price of manufactures. These workers instead enter the other three sectors, but predominantly the service sector. Because the service sector is nontraded, the change in the entry pattern leads to a decrease in the price of services and ultimately the change in world prices is responsible for a 15 percent decline in the price of services. As a consequence, even though the decrease in world prices decreases real prices in the economy by 10 percent in the long

run, both manufacturing and service sector workers are hurt because the prices of the goods they produce fall by even more.

In the long run, the average worker in the natural resource sector is approximately 8 percent better off as compared to the average worker in the natural resource sector had prices remained unchanged. However, the story is different in the other sectors. Workers in the agriculture, services and manufacturing sectors are, respectively, 3, 8 and 9 percent worse off. Taken together, the world price changes are responsible for an aggregate welfare decline of 7.5 percent in the short run and 5.5 percent in the long run. The reason for the decline in aggregate welfare is that workers are not so easily swayed to move into the positively impacted sectors due to the distribution of idiosyncratic sector-entry-cost draws. In other words, Canadian workers are hurt on average by the change in world prices because the price changes do not lead to a large influx of workers into the sectors that experience positive shocks.

The research in this chapter represents an important step in addressing the types of concerns raised in the public debate regarding increased global integration. These concerns, such as, the effect of globalization on unemployment and labor-force participation among others addressed in this paper, have yet to receive adequate treatment in the literature. This chapter focuses on the differential impact of global changes on workers, but incorporating other factors of production may provide more important insights into the effect of global changes on a small open economy. Future research may also incorporate data on trade flows or a model of the global economy in order to identify precisely which changes in the global economy have been the ultimate cause of changing world prices.

Chapter 3

A Pareto-Imroving Minimum Wage

This chapter shows that a graduated minimum wage, in contrast to a constant minimum wage, can provide a strict Pareto improvement over what can be achieved with an optimal income tax. The reason is that a graduated minimum wage requires high-productivity workers to work more to earn the same income as low-productivity workers, which makes it more difficult for the former to mimic the latter. In effect, a graduated minimum wage allows the low-productivity workers to benefit from second-degree price discrimination which increases their income.[1]

3.1 Introduction

In a competitive economy, a constant minimum wage cannot be used to make everyone better off and thereby provide a strict Pareto improvement over what can be achieved by an optimal income tax alone. Indeed, as shown by Allen (1987) and Guesnerie and Roberts (1987), a binding constant minimum wage would be counterproductive as it would reduce the hours of low-productivity workers. The reason is that the government could have attained the same lower working hours by an appropriately designed income tax alone, but has refrained from

[1]This chapter was coauthored with Leif Danziger. A version of this chapter has been accepted for publication in Economica.

doing so. In other words, the power to set a constant minimum wage does not provide the government with an extra instrument to affect the feasible consumption-work bundles.

Subsequent papers have tried to justify the minimum wage by modifying the competitive framework and changing the objective to be to increase social welfare rather than to provide a strict Pareto improvement. Thus, if the government's social welfare function expresses a sufficiently strong taste for redistribution toward the less productive workers, combining a constant minimum wage with an optimal income tax may lead to an increase in social welfare under various circumstances. For example, if it forces some low-productivity workers to become unemployed (Marceau and Boadway, 1994); if it is combined with a welfare policy that obliges workers to accept job offers (Boadway and Cuff, 2001); if it prevents workers from signalling their earning ability (Blumkin and Sadka, 2005); or if workers with the highest disutilities of work are laid off first (Lee and Saez, 2012).

In the absence of optimal income taxation, a constant minimum wage may increase social welfare both in a competitive environment (Danziger, 2009a; and Lee and Saez, 2012)[2] and in a monopsonistic environment (Robinson, 1933).[3] Since search frictions and informational asymmetries may lead to monopsonistic features, a constant minimum wage may be rationalized with other types of labor markets if income taxation is not optimal (Rebitzer and Taylor, 1995; Bhaskar and To, 1999; Cahuc et al., 2001; Bhaskar et al., 2002; Manning, 2003; Flinn, 2006; Hungerbühler and Lehmann, 2009; and Basu et al., 2010). Additionally, politics, the power of unions, and cultural factors may affect the setting of the minimum wage (Sobel, 1999; Belot et al., 2007; Boeri and Burda, 2009; Brown, 2009; Checchi and Garcia-Penalosa, 2010; Aghion et al., 2011; and Boeri, 2012).

In the current chapter we pursue a different path and provide a novel justification for the minimum wage without modifying the competitive framework or changing the objective function. Our approach is designed to be as close as possible to the standard competitive

[2]See also Stewart and Swaffield (2008), De Fraja (1999), and Strobl and Walsh (2011).

[3]However, with optimal income taxation, monopsony cannot justify a constant minimum wage (Cahuc and Laroque, in press).

model with two types of workers and optimal income taxes (Stiglitz, 1982). As we will show, our minimum wage design can provide a strict Pareto improvement of the second-best optimum achievable with taxation alone.

There is an obvious dissonance between insisting on the minimum wage being constant while allowing for a nonlinear income tax schedule. Thus, in the spirit of the optimal income tax literature, this paper analyzes a graduated (i.e., nonlinear) minimum wage that depends on the total employment of a firm's minimum-wage workers. We show that a graduated minimum wage severs the tight connection between the wage and a workers' marginal product, thereby providing a tool for compelling firms to hire low-productivity workers at a minimum wage that exceeds their marginal product.

A graduated minimum wage can be designed so that a high-productivity worker needs to work more hours in order to earn the same income as a low-productivity worker. This makes it less attractive for high-productivity workers to mimic the low-productivity workers by earning their income. Consequently, it is easier to distinguish between workers with different productivities, and therefore also to redistribute toward the less productive workers. We prove that even if the economy is initially in a second-best optimum with an optimal income tax, both with homogenous and heterogenous firms the introduction of a graduated minimum wage can benefit all individuals and thus constitute a strict Pareto improvement.

As pointed out by Guesnerie and Roberts (1987), there is an apparent informational inconsistency in assuming that while the government cannot observe the workers' wages (and hence condition taxes on wages), it can nevertheless enforce a minimum wage. However, Lee and Saez (2012, p. 746) argue that even though it may not be possible for the government to observe wages, enforcement of a minimum wage can be brought about by a system of audits triggered by worker complaints and hefty penalties for noncompliant firms.[4] Furthermore, as we show that there exists a graduated minimum wage schedule for which the intended minimum wage is less than the minimum wage at all but the intended working hours, a

[4] Alternatively, it may be the case that a governmental enforcement agency can observe workers' wages, but that political or institutional constraints limit taxes to depend on only income.

worker's complaint of being paid less than the specified minimum wage (given the firm's total employment of low-productivity workers) is equivalent to a complaint that the firm's total employment of low-productivity workers is improper (given the minimum wage).

Nevertheless, enforcement of a graduated minimum wage is likely to be more complicated and require more resources than enforcement of a constant minimum wage. However, other government programs also impose hour requirements that necessitate keeping track of the number of working hours and therefore encounter similar problems. For example, in the US, the Work Opportunity Tax Credit allows firms to claim tax credit equal to 25% of an eligible worker's first-year earnings if the individual works at least 120 hours, and equal to 40% of an eligible worker's first-year earnings if the individual works at least 400 hours. Similarly, the Welfare-to-Work Tax Credit provides firms with a tax credit of 35% of an eligible worker's first-year earnings and 50% of an eligible worker's second-year earnings if the individual works at least 400 hours.[5]

If the graduated minimum wage is differentiable, we show that a firm's marginal labor cost is less than the minimum wage. Effectively, therefore, firms are facing a downward-sloping labor supply of low-productivity workers so that these workers are employed more than they would be at a constant minimum wage equal to their marginal product. In other words, the graduated minimum wage endows the workers with monopoly power that allows the practice of second-degree price discrimination which increases their income.

3.2 The Standard Competitive Model with Taxes

We consider an economy with a single consumption good. There is a continuum of low-productivity workers of measure $n_1 > 0$ and a continuum of high-productivity workers of measure $n_2 > 0$. The workers share the same preferences and their utility function is $U(c, h)$, where $c \geq 0$ is the worker's consumption and $h \geq 0$ is the working hours. The utility function

[5]The issue of noncompliance is studied in Ashenfelter and Smith (1979), Weil (2005), and Danziger (2009b, 2010).

satisfies $U_c > 0$, $U_{cc} < 0$, $U_h < 0$, $U_{hh} \leq 0$, and the agent monotonicity property, i.e., hU_h/U_c is decreasing in h.

The economy also includes a unit continuum of profit-maximizing homogenous firms. A firm's production function is $f(\ell_1, \ell_2)$, where ℓ_1 and ℓ_2 are the total working hours in the firm of the low- and high-productivity workers, respectively. The production functions satisfy $f(0,0) = 0$ and exhibit decreasing returns to scale, a decreasing marginal product of the low-productivity workers, and an imperfect substitutability between the two types of workers (i.e., the isoquants are not straight lines). Due to a bias in the production functions in favor of the high-productivity workers and/or a large ratio of low- to high-productivity workers, in all relevant allocations a low-productivity worker's marginal product is less than that of a high-productivity worker. This justifies our classification of workers as having either low or high productivity.

The low- and high-productivity workers' competitive wages are w_1 and w_2, respectively, where $w_1 < w_2$. Assuming that working hours are positive for both types of workers, the decreasing returns to scale imply that profits are positive in a competitive equilibrium, i.e., $f(\ell_1, \ell_2) - w_1\ell_1 - w_2\ell_2 > 0$. The government taxes all profits away.[6] Also, we let (c_1, h_1) and (c_2, h_2) denote the consumption and working hours of the low- and high-productivity workers, respectively, where $h_1 \equiv \ell_1/n_1$ and $h_2 \equiv \ell_2/n_2$.

In addition to workers and firms, there is a government that wants to transfer consumption from the high- to the low-productivity workers (i.e., $c_1 - w_1h_1 > c_2 - w_2h_2$). As is standard in the tax literature, taxation is based on incomes and cannot depend directly on wages. Since this rules out a first-best allocation in which lump-sum taxation is based on productivity levels, the government chooses a second-best allocation by determining a non-linear income tax that transfers consumption from the high- to the low-productivity workers. The government's choice has to satisfy the incentive-compatibility constraints that workers of one type cannot make themselves better off by choosing their working hours so that they

[6] We later discuss how the assumptions can be modified so that either profits are zero or belong to firm owners who are not fully taxed.

earn the same income as workers of the other type. To write this formally, let a circumflex over a variable denote its value in a second-best optimum with nonlinear income taxes and redistribution toward the low-productivity workers, but without a minimum wage. Then the low- and high-productivity workers' incentive-compatibility constraints are

$$U(\hat{c}_1, \hat{h}_1) \geq U\left(\hat{c}_2, \frac{\hat{w}_2\hat{h}_2}{\hat{w}_1}\right), \tag{3.1}$$

$$U(\hat{c}_2, \hat{h}_2) \geq U\left(\hat{c}_1, \frac{\hat{w}_1\hat{h}_1}{\hat{w}_2}\right), \tag{3.2}$$

where $\hat{w}_2\hat{h}_2/\hat{w}_1$ (resp. $\hat{w}_1\hat{h}_1/\hat{w}_2$) are the hours that a low-productivity (resp. high-productivity) worker would have to work in order to reach a high-productivity (resp. low-productivity) worker's income. The agent monotonicity property implies that the incentive-compatibility constraint of the low-productivity workers (3.1) is slack in a second-best optimum, while that of the high-productivity workers (3.2) binds (Salanié, 2003).

The government's choice must also satisfy the resource constraint

$$f(\hat{\ell}_1, \hat{\ell}_2) - n_1\hat{c}_1 - n_2\hat{c}_2 = R, \tag{3.3}$$

where R is the government's exogenous spending need. In sum, in a second-best redistributive optimum the government maximizes $U(c_1, h_1)$ subject to a given level of $U(c_2, h_2)$ and the binding constraints (3.2) and (3.3).

3.3 A Graduated Minimum Wage

Suppose that the low-productivity workers must be paid according to a graduated minimum wage schedule $m(\ell_1)$ which would make the lowest wage a firm can pay a function of the total working hours of all the low-productivity workers employed in the firm. Since any potential minimum wage must be the lowest wage in the economy (even if never chosen), $m(\ell_1)$ cannot

exceed a high-productivity worker's wage for any ℓ_1; that is, $m(\ell_1) \leq w_2$ for any ℓ_1. Letting an asterisk indicate the value of a variable in an equilibrium with a graduated minimum wage, the low- and high-productivity workers' incentive-compatibility constraints become

$$U(c_1^*, h_1^*) \geq U\left(c_2^*, \frac{w_2^* h_2^*}{w_1^*}\right), \tag{3.4}$$

$$U(c_2^*, h_2^*) \geq U\left(c_1^*, \frac{w_1^* h_1^*}{w_2^*}\right). \tag{3.5}$$

Just as the income-tax function offers the workers a choice between different consumption-work bundles, the graduated minimum wage offers the firms a choice between different minimum wage-hour bundles. With the optimal income tax one need only be concerned with determining the intended consumption-work bundles for the two types of workers (because all other bundles can be made arbitrarily unattractive). However, with the optimal graduated minimum wage one need be concerned not only with determining the intended minimum wage-hours bundle, but also with making sure that a firm's profit with that bundle is at least as high as the profit that can be attained with any other minimum wage-hours bundle. That is, a firm's choice of working hours of low- and high-productivity workers must satisfy the minimum wage constraint

$$f(\ell_1^*, \ell_2^*) - m(\ell_1^*)\ell_1^* - w_2^*\ell_2^* \geq f(\ell_1, \ell_2) - m(\ell_1)\ell_1 - w_2^*\ell_2 \qquad \forall(\ell_1, \ell_2). \tag{3.6}$$

Since a firm has the option of not producing by setting $\ell_1 = \ell_2 = 0$ and hence earning zero profit, there exist working hours (ℓ_1^*, ℓ_2^*) for which a firm's profit is nonnegative and the minimum wage constraint is satisfied.

The resource constraint becomes

$$f(\ell_1^*, \ell_2^*) - n_1 c_1^* - n_2 c_2^* = R, \tag{3.7}$$

and in an equilibrium the government's choice of a nonlinear income tax and a graduated minimum wage must satisfy the low- and the high-productivity workers' incentive-compatibility constraints (3.4) and (3.5), the firms' minimum wage constraint (3.6), and the resource constraint (3.7).

We will now show that even if the economy is initially in a second-best optimum with a nonlinear income tax that redistributes toward the less productive workers, the introduction of a graduated minimum wage can achieve a strict Pareto improvement:

Proposition 3.3.1 *With homogenous firms a graduated minimum wage can provide a strict Pareto improvement of a second-best optimum with nonlinear income taxation that redistributes consumption toward the less productive workers.*

Proof: In three steps we construct a graduated minimum wage that increases the utility of all workers in equilibrium.

Step 1: *A Graduated Minimum Wage that Leaves the Utility of All Workers Unchanged*

Suppose that the graduated minimum wage be given by

$$m(\ell_1) = \begin{cases} \hat{w}_1 + \alpha & \text{for } \ell_1 = \hat{\ell}_1, \\ \hat{w}_2 & \text{for } \ell_1 \neq \hat{\ell}_1, \end{cases}$$

where $0 < \alpha < \hat{w}_2 - \hat{w}_1$. That is, the wage is α higher than $\hat{w}_1$ at the same total hours of low-productivity workers as in the second-best allocation with only nonlinear income taxes, and $\hat{w}_2 - \hat{w}_1$ higher than $\hat{w}_1$ at all other hours.

Let the intended consumption-work bundle for the low-productivity workers remain unchanged, which requires that their tax payment increases by $\alpha\hat{h}_1$. Also, let the intended consumption-work bundle for the high-productivity workers remain unchanged, which requires that the income tax when earning $\hat{w}_2\hat{h}_2$ be unchanged. The low-productivity workers' incentive-compatibility constraint (3.4) will then be the same as (3.1) with $\hat{w}_1+\alpha$ substituted for $\hat{w}_1$, and it will remain slack if α is small enough; that is,

$$U(\hat{c}_1, \hat{h}_1) > U\left(\hat{c}_2, \frac{\hat{w}_2\hat{h}_2}{\hat{w}_1 + \alpha}\right). \tag{3.8}$$

The high-productivity workers have to work more in order to earn the income of the low-productivity workers. Therefore, their incentive-compatibility constraint (3.5) will be the same as (3.2) with $\hat{w}_1 + \alpha$ substituted for $\hat{w}_1$. As the right-hand-side of (3.2) decreases, their incentive-compatibility constraint is loosened and will be slack; that is

$$U(\hat{c}_2, \hat{h}_2) > U\left[\hat{c}_1, \frac{(\hat{w}_1 + \alpha)\hat{h}_1}{\hat{w}_2}\right]. \tag{3.9}$$

If it were the case that $\alpha = 0$ so that $m(\hat{\ell}_1) = \hat{w}_1$, but $m(\ell_1) = \hat{w}_2$ for $\ell_1 \neq \hat{\ell}_1$, then the left-hand-side of the minimum wage constraint (3.6) would equal a firm's profit in the absence of a minimum wage, while the right-hand-side for any $\ell_1 \neq \hat{\ell}_1$ would be less since the hourly wage of low-productivity workers have increased from $\hat{w}_1$ to $\hat{w}_2$ for any $\ell_1 \neq \hat{\ell}_1$. The constraint would therefore hold with a strict inequality for all (ℓ_1, ℓ_2) with $\ell_1 \neq \hat{\ell}_1$. In reality, though, the intended minimum wage exceeds $\hat{w}_1$ by $\alpha > 0$ so that a firm's profit at $(\hat{\ell}_1, \hat{\ell}_2)$ decreases by $\alpha\hat{\ell}_1$. However, for a sufficiently small α and any given $\ell_1 > 0$ we have that $\alpha\hat{\ell}_1 < (\hat{w}_2 - \hat{w}_1)\ell_1$ so that for small α's the left-hand-side of the minimum wage constraint (3.6) is less than the right-hand-side for all (ℓ_1, ℓ_2) with $\ell_1 \neq \hat{\ell}_1$. This implies that for a sufficiently small α the minimum wage constraint becomes slack and a firm's profit is positive; that is

$$f(\hat{\ell}_1, \hat{\ell}_2) - (\hat{w}_1 + \alpha)\hat{\ell}_1 - \hat{w}_2\hat{\ell}_2 > f(\ell_1, \ell_2) - \hat{w}_2\ell_1 - \hat{w}_2\ell_2 \qquad \forall(\ell_1, \ell_2). \tag{3.10}$$

In actuality, therefore, the graduated minimum wage compels the firms to pay low-productivity workers more than the competitive wage while preventing them from shrinking the labor input of these workers.

Since the workers' consumption and working hours do not change, the resource constraint will not be affected. As a result, the graduated minimum wage schedule detailed above leads to an equilibrium in which the incentive-compatibility constraints of both low- and high-productivity workers as well as the minimum wage constraint are slack. Nonetheless, the consumption and working hours of each type of worker, and hence their utilities, will be the same as in a second-best optimum with nonlinear income taxes and redistribution toward the low-productivity workers.

Step 2: *A Graduated Minimum Wage that Increases the Utility of Low-Productivity Workers*

In a second-best optimum with nonlinear income taxes and redistribution toward the low-productivity workers, their marginal rate of substitution between consumption and leisure will be less than their marginal product (see Appendix). Therefore, the low-productivity workers would benefit if their working hours were increased and they would get to consume all the additional production. One implication is that there exists an increase in each low-productivity worker's hours $\gamma > 0$ such that $U(\hat{c}_1 + \delta, \hat{h}_1 + \gamma) > U(\hat{c}_1, \hat{h}_1)$, where $\delta \equiv [f(\hat{\ell}_1 + n_1\gamma, \hat{\ell}_2) - f(\hat{\ell}_1, \hat{\ell}_2)]/n_1$. Here, $n_1\gamma$ is the increase in the total working hours of low-productivity workers in a firm (because there is a unit continuum of firms) and $n_1\delta$ is the corresponding increase in the output.

We now let the graduated minimum wage for $\hat{\ell}_1 + n_1\gamma$ working hours of low-productivity workers in a firm be equal to $[(\hat{w}_1 + \alpha)\hat{h}_1 + \delta]/(\hat{h}_1 + \gamma)$, and the income tax when earning $m(\hat{\ell}_1 + n_1\gamma)(\hat{h}_1 + \gamma) = (\hat{w}_1 + \alpha)\hat{h}_1 + \delta$ be such that the low-productivity workers' consumption is $\hat{c}_1 + \delta$. We also let the high-productivity workers' income tax be adjusted such that if their wage changes to $\tilde{w}_2 \equiv \partial f(\hat{\ell}_1 + n_1\gamma, \hat{\ell}_2)/\partial\hat{\ell}_2$, then their consumption and working hours remain unchanged. Finally, we let the graduated minimum wage be equal to $\tilde{w}_2$ for $\ell_1 \neq \hat{\ell}_1 + n_1\gamma$.

The low- and high-productivity workers' incentive-compatibility constraints are

$$U(\hat{c}_1 + \delta, \hat{h}_1 + \gamma) \geq U\left[\hat{c}_2, \frac{\tilde{w}_2 \hat{h}_2}{m(\hat{\ell}_1 + n_1\gamma)}\right], \tag{3.11}$$

$$U(\hat{c}_2, \hat{h}_2) \geq U\left[\hat{c}_1 + \delta, \frac{(\hat{w}_1 + \alpha)\hat{h}_1 + \delta}{\tilde{w}_2}\right], \tag{3.12}$$

and the minimum wage constraint is

$$f(\hat{\ell}_1 + n_1\gamma, \hat{\ell}_2) - m(\hat{\ell}_1 + n_1\gamma)\hat{\ell}_1 - \tilde{w}_2\hat{\ell}_2 \geq f(\ell_1, \ell_2) - \tilde{w}_2\ell_1 - \tilde{w}_2\ell_2 \qquad \forall(\ell_1, \ell_2). \tag{3.13}$$

Since the constraints (3.11), (3.12), and (3.13) would be identical to (3.8), (3.9), and (3.10) for $\gamma = \delta = 0$, and the latter are slack, there exist positive values of γ for which (3.11), (3.12), and (3.13) are also slack. The resource constraint would also be satisfied. Hence, there is a graduated minimum wage leading to an equilibrium in which the low-productivity workers's utility increases while the high-productivity workers' utility remains unchanged.

Step 3: *A Graduated Minimum Wage that Increases the Utility of All Workers*

Consider a value of γ for which (3.11), (3.12), and (3.13) are slack. The fact that the low-productivity workers gain and the high-productivity workers are indifferent implies that there exists a small $\epsilon > 0$ such that if, for the same intended minimum wage as in step 2, the low-productivity workers' consumption is reduced from $\hat{c}_1 + \delta$ to $\hat{c}_1 + (1 - \epsilon)\delta$ and the high-productivity workers' consumption is increased from $\hat{c}_2$ to $\hat{c}_2 + \epsilon\delta n_1/n_2$, then:

- both the low-productivity workers' utility $U[\hat{c}_1 + (1 - \epsilon)\delta, \hat{h}_1 + \gamma]$ and the high-productivity workers' utility $U(\hat{c}_2 + \epsilon\delta n_1/n_2, \hat{h}_2)$ will be higher than in the second-best optimum with nonlinear income taxes and redistribution toward the low-productivity workers;

- both the low- and the high-productivity workers' incentive-compatibility constraints will be slack, i.e.,

$$\begin{aligned} U[\hat{c}_1 + (1-\epsilon)\delta, \hat{h}_1 + \gamma] &> U\left[\hat{c}_2 + \frac{\epsilon\delta n_1}{n_2}, \frac{\tilde{w}_2\hat{h}_2}{m(\hat{\ell}_1 + n_1\gamma)}\right], \\ U\left(\hat{c}_2 + \frac{\epsilon\delta n_1}{n_2}, \hat{h}_2\right) &> U\left[\hat{c}_1 + (1-\epsilon)\delta, \frac{(\hat{w}_1 + \alpha)\hat{h}_1 + \delta}{\tilde{w}_2}\right]; \end{aligned}$$

- the minimum wage constraint will be slack, i.e.,

$$f(\hat{\ell}_1 + n_1\gamma, \hat{\ell}_2) - m(\hat{\ell}_1 + n_1\gamma)(\hat{\ell}_1 + n_1\gamma) - \tilde{w}_2\hat{\ell}_2 > f(\ell_1, \ell_2) - \tilde{w}_2\ell_1 - \tilde{w}_2\ell_2 \quad \forall(\ell_1, \ell_2).$$

Since the resource constraint continues to hold, we have proved that there exists an equilibrium with the graduated minimum wage

$$m(\ell_1) = \begin{cases} w_1^* & \text{for } \ell_1 = \ell_1^*, \\ \tilde{w}_2 & \text{for } \ell_1 \neq \ell_1^*, \end{cases}$$

where $w_1^* = [(\hat{w}_1 + \alpha)\hat{h}_1 + \delta]/(\hat{h}_1 + \gamma)$. In this equilibrium, $c_1^* = \hat{c}_1 + (1-\epsilon)\delta$; $h_1^* = \hat{h}_1 + \gamma$; $c_2^* = \hat{c}_2 + \epsilon\delta n_1/n_2$; $h_2^* = \hat{h}_2$; and $w_2^* = \tilde{w}_2$. All workers obtain a higher utility than in the second-best optimum with only nonlinear income taxes and redistribution toward the low-productivity workers. This completes the proof. □

As the construction of the proof of Proposition 3.3.1 shows, a strict Pareto improvement is possible because a graduated minimum wage is more efficacious than income taxes in directing resources to the low-productivity workers. Specifically, the graduated minimum wage funnels more of the firms' revenues to the low-productivity workers, thereby raising their pretax income without affecting that of the high-productivity workers (see step 1 of the proof). High-productivity workers will then have to work more in order to earn the same pretax income as low-productivity workers, and this mitigates the high-productivity workers' incentive-compatibility constraint. Consequently, the graduated minimum wage

makes it easier for the government to distinguish between the low- and high-productivity workers, which facilitates a strict Pareto improvement.

The graduation of the minimum wage is essential for achieving a strict Pareto improvement, as Allen (1987) and Guesnerie and Roberts (1987) have shown that a constant minimum wage cannot do so. This is because, on the one hand, a constant minimum wage that does not exceed the competitive wage is ineffectual, while, on the other, a constant minimum wage that does exceed the competitive wage will reduce the working hours of the low-productivity workers. However, the same lower working hours could also have been obtained with a nonlinear income tax alone. The fact that the government has chosen not to do so reveals that it would be deleterious.

Our model assumes that the measure of firms is given and hence that there is no free entry. Together with the decreasing returns to scale and the absence of a fixed cost of production, this implies that firms earn positive profits in a second-best optimum with nonlinear taxation. These profits are fully taxed by the government, and when moving to an equilibrium with a graduated minimum wage, the graduated minimum wage is instrumental in transferring part of the profits to the low-productivity workers while the government continues to fully tax the remaining profits. Indeed, a key ingredient in the Pareto-improving mechanism is that firms can pay the low-productivity workers more than their marginal product without losing money.

If the assumption that the measure of firms is given and hence that firms earn positive profits is a concern, one could alternatively assume that the measure of firms is not given but that production involves a fixed cost. With decreasing returns to scale, firms would enter until the profits are driven down to zero in a second-best optimum with nonlinear taxation, and the economy's resource constraint would be modified by subtracting the firms' fixed costs on the left-hand side of the constraint. The government could introduce a graduated minimum wage in the same way as in the proof of Proposition 3.3.1. This would then also involve subsidizing any active firm with an amount equal to how much extra the firm

is obliged to pay the low-productivity workers due to the graduated minimum wage, with the government financing the subsidies by increased taxation in the same way as it would finance an increase in its exogenous spending need for any other reason.[7] This would leave the measure of firms unchanged as their profits would remain equal to zero. Accordingly, with free entry of firms and a fixed cost of production, introducing a graduated minimum wage together with a subsidy to active firms is equivalent to introducing a graduated minimum wage when the measure of firms is given, there is no fixed cost of production, and the government's spending need is raised by an amount equal to the subsidies. It follows that Proposition 3.3.1 would hold also in a case with free entry of firms, fixed cost of production, and subsidies to active firms.[8]

If one keeps the assumption of a given measure of firms earning positive profits, the full taxation of the profits is important since there is then no need to worry about firm owners. However, instead of full taxation of profits, one could assume that only profits above the normal level are taxed away. The resource constraint for the economy would then be modified by adding the normal profit to the government's exogenous spending need, and Proposition 3.3.1 would state that a graduated minimum wage can provide a strict Pareto improvement for the workers while leaving the firm owners no worse off. By applying a similar argument as in step 3 of the proof of Proposition 3.3.1, the proposition could be generalized to state that there exists a graduated minimum wage which can provide a strict Pareto improvement for both the workers and the firm owners.

[7] Government subsidy to firms can take various forms including active labor market programs (such as the Work Opportunity Tax Credit and the Welfare-to-Work Tax Credit mentioned in the introduction), guaranteed loans, start-up grants, farm subsidies, R&D support, and preferential tax treatments.

[8] One could also avoid positive profits and have Proposition 1 hold by assuming constant-returns-to-scale and no fixed cost, together with a tax-financed government subsidy to active firms equaling the extra amount the graduated minimum wage forces the firms to pay the low-productivity workers.

3.4 A Continuous Piecewise Differentiable Graduated Minimum Wage

In our formulation until now, the graduated minimum wage is not continuous at the intended working hours of the low-productivity workers (since it equals $[(\hat{w}_1 + \alpha)\hat{h}_1 + \delta]/(\hat{h}_1 + \gamma)$ for the intended working hours and equals $\tilde{w}_2$ for other working hours). To shed further light on the source of the strict Pareto improvement, assume now that the minimum wage schedule, which we still write as $m(\ell_1)$, is locally differentiable at the intended working hours of the low-productivity workers. To maximize its profit, a firm would then set the low-productivity workers' hours such that the marginal labor cost, $\partial[m(\ell_1)\ell_1]/\partial\ell_1 = m(\ell_1) + m'(\ell_1)\ell_1$, equals their marginal product. Accordingly, if the minimum wage constraint is satisfied, profit maximization implies that[9]

$$m\left(\ell_1\right) + m'\left(\ell_1\right)\ell_1 = f_{\ell_1}. \tag{3.14}$$

In an equilibrium with a Pareto-improving graduated minimum wage, each firm employs low-productivity workers for more hours than it would choose if the wage were constant at $\hat{w}_1$ and the firm could freely decide the hours. Accordingly, the minimum wage exceeds a low-productivity worker's marginal product, i.e., $m\left(\ell_1\right) > f_{\ell_1}$. By (3.14), this implies that the marginal labor cost of low-productivity workers, $m(\ell_1) + m'(\ell_1)\ell_1$, is less than the minimum wage, $m(\ell_1)$, and hence that the marginal minimum wage, $m'\left(\ell_1\right)$, is negative. That is, a strict Pareto improvement would be achieved by a locally differentiable minimum wage schedule that is downward sloping at the intended hours. In effect, the graduated minimum wage transforms the low-productivity workers into second-degree price discriminating monopolists that confront firms with a downward-sloping labor supply at the intended hours. The result is that the low-productivity workers are able to extract some of the firms' rev-

[9] The minimum wage constraint ensures that the second-order condition for profit maximization holds, i.e., that $2m'\left(\ell_1\right) + m''\left(\ell_1\right)\ell_1 > f_{\ell_1\ell_1}$.

enues and that they earn more than they would if the firms were permitted to pay the same minimum wage for fewer working hours.

We can now establish:

Proposition 3.4.1 *There exists a continuous piecewise differentiable graduated minimum wage schedule and income taxes that can implement the same strict Pareto improvement as the non-continuous graduated minimum wage constructed in the proof of Proposition 1.*

Proof: Let the intended minimum wage $m(\ell_1^*)$ and the income taxes be the same as in the proof of Proposition 1. Further, let the graduated minimum wage schedule be given by

$$m(\ell_1) = \begin{cases} f_{\ell_1^*} + \dfrac{\beta}{\ell_1} & \text{for } \ell_1 > \dfrac{\beta}{\tilde{w}_2 - f_{\ell_1^*}}, \\ \tilde{w}_2 & \text{for } \ell_1 \leq \dfrac{\beta}{\tilde{w}_2 - f_{\ell_1^*}}, \end{cases}$$

where $\beta \equiv (w_1^* - f_{\ell_1^*})\ell_1^*$ and $f_{\ell_1^*}$ is evaluated at (ℓ_1^*, ℓ_2^*). Note that the minimum wage is always continuous and that the marginal labor cost is $f_{\ell_1^*}$ for $\ell_1 > \beta/(\tilde{w}_2 - f_{\ell_1^*})$ and $\tilde{w}_2$ for $\ell_1 \leq \beta/(\tilde{w}_2 - f_{\ell_1^*})$. Thus, profit maximization implies that a firm sets the low-productivity workers' hours to ℓ_1^* and that the corresponding (intended) graduated minimum wage is w_1^*.

Since income taxes are the same as before, it is immediate that the low- and high-productivity workers' incentive-compatibility constraints and the minimum wage constraint will be slack. It is also clear that the resource constraint still holds. Hence, the same Pareto-improving equilibrium as constructed in the proof of Proposition 1 is achieved. □

The graduated minimum wage used in the proof of Proposition 3.4.1 is, in the relevant range, a hyperbolic function of the low-productivity workers' hours. This implies that the marginal labor cost is a constant that equals $f_{\ell_1^*}$ independently of the value of β. Therefore, the hours are the same as they would be with a competitive wage w_1^*, but the low-productivity workers nevertheless succeed in increasing their total pay by an additional β.

3.5 Heterogeneous Firms

We now generalize the model by allowing for heterogenous firms. Thus, suppose that there are two categories of profit-maximizing firms, with x_j being the measure of firms in category $j \in \{1, 2\}$ and $x_1 + x_2 = 1$. The production function of a firm in category j is $f_j(\ell_{1j}, \ell_{2j})$, where ℓ_{1j} and ℓ_{2j} are the total hours of labor input of the low- and high-productivity workers, respectively. The production functions satisfy $f_j(0, 0) = 0$ and exhibit decreasing returns to scale, a decreasing marginal product of the low-productivity workers, and an imperfect substitutability between the two types of workers. Each worker can only work in one firm.

In the competitive framework with nonlinear income taxes but no graduated minimum wage, we continue to denote the low- and high-productivity workers' competitive wages by w_1 and w_2, respectively, and the high-productivity workers' consumption and working hours by (c_2, h_2). We let $n_{1j} > 0$, where $n_{11}+n_{12} = n_1$, denote the measure of low-productivity workers who work for a firm in category j, and c_{1j} and $h_{1j} \equiv \ell_{1j}x_j/n_{1j}$ denote the consumption and working hours of low-productivity workers who work for a firm in category j. The working hours are assumed positive for both types of workers in both categories of firms. It follows that in equilibrium all low-productivity workers will obtain the same consumption-work bundle no matter where they work, i.e., $(\hat{c}_{11}, \hat{h}_{11}) = (\hat{c}_{12}, \hat{h}_{12}) = (\hat{c}_1, \hat{h}_1)$. Furthermore, the decreasing returns to scale imply that in a second-best optimum with nonlinear income taxes, the profits are positive for both firm categories.

The low- and high productivity workers' incentive-compatibility constraints are again given by (3.1) and (3.2), with only the latter being binding, while the resource constraint is

$$x_1 f_1(\hat{\ell}_{11}, \hat{\ell}_{21}) + x_2 f_2(\hat{\ell}_{12}, \hat{\ell}_{22}) - n_1\hat{c}_1 - n_2\hat{c}_2 = R. \tag{3.15}$$

In a second-best optimum with redistribution favoring the low-productivity workers, the government chooses the nonlinear income tax to maximize $U(c_1, h_1)$ for a given level of $U(c_2, h_2)$ subject to the binding constraints (3.2) and (3.15).

After the introduction of the graduated minimum wage, with heterogenous firms the incentive-compatibility constraints that the low- and high-productivity workers cannot improve their utility by choosing working hours that yield the income of the other type of worker are generalized to

$$U(c_{1j}^*, h_{1j}^*) \geq U\left[c_2^*, \frac{w_2^* h_2^*}{m(\ell_{1j}^*)}\right] \quad \text{for } j \in \{1,2\}, \tag{3.16}$$

$$U(c_2^*, h_2^*) \geq U\left[c_{1j}^*, \frac{m(\ell_{1j}^*) h_{1j}^*}{w_2^*}\right] \text{ for } j \in \{1,2\}. \tag{3.17}$$

Further, the firms' choices of working hours for the low- and high-productivity workers must satisfy the minimum wage constraint

$$f_j(\ell_{1j}^*, \ell_{2j}^*) - m(\ell_{1j}^*)\ell_{1j}^* - w_2^*\ell_{2j}^* \geq f_j(\ell_{1j}, \ell_{2j}) - m(\ell_{1j})\ell_{1j} - w_2^*\ell_{2j} \qquad \forall(\ell_1, \ell_2) \text{ for } j \in \{1,2\}. \tag{3.18}$$

Since it is possible for a firm to set $\ell_{1j} = \ell_{2j} = 0$ and obtain zero profit, for $j \in \{1,2\}$ there exist labor inputs $(\ell_{1j}^*, \ell_{2j}^*)$ which yield a nonnegative profit and satisfy the minimum wage constraint.

In the presence of heterogenous firms, the graduated minimum wage introduces an additional incentive-compatibility constraint for the low-productivity workers. The intended consumption-work bundles of the low-productivity workers employed in the different categories of firms must be such that those employed in one category of firms cannot make themselves better off by earning the same income as those employed in the other category.[10] Thus, the additional incentive-compatibility constraint for low-productivity workers is

$$U(c_{1j}^*, h_{1j}^*) \geq U\left[c_{1j'}^*, \frac{m(\ell_{1j'}^*) h_{1j'}^*}{m(\ell_{1j}^*)}\right] \text{ for } j \neq j'. \tag{3.19}$$

[10] All low-productivity workers prefer to work for the firms that pay the higher minimum wage. However, while some low-productivity workers have the good luck to be employed in those firms, others have no choice but to resign themselves to work for the lower-paying firms.

The resource constraint is

$$x_1 f_1(\ell_{11}^*, \ell_{21}^*) + x_2 f_2(\ell_{12}^*, \ell_{22}^*) - n_{11} c_{11}^* - n_{12} c_{12}^* - n_2 c_2^* = R. \tag{3.20}$$

The government's choice of a nonlinear income tax and a graduated minimum wage must satisfy the low-productivity workers' incentive-compatibility constraints (3.16) and (3.19), the high-productivity workers' incentive-compatibility constraint (3.17), the firms' minimum wage constraint (3.18), and the resource constraint (3.20).

We now extend Proposition 1 to the case of heterogenous firms:

Proposition 3.5.1 *In the case of heterogenous firms there also exists a graduated minimum wage which can provide a strict Pareto improvement of a second-best optimum with nonlinear income taxation that redistributes consumption toward the less productive workers.*

Proof: Following the same logic as in the proof of Proposition 3.3.1 , we use three steps to construct an equilibrium with a graduated minimum wage that increases the utility of all workers.

Step 1: Suppose the graduated minimum wage schedule is

$$m(\ell_{1j}) = \begin{cases} \hat{w}_1 + \sigma & \text{for } \ell_{1j} \in \{\hat{\ell}_{11}, \hat{\ell}_{12}\}, \\ \hat{w}_2 & \text{for } \ell_{1j} \notin \{\hat{\ell}_{11}, \hat{\ell}_{12}\}, \end{cases}$$

where $0 < \sigma < \hat{w}_2 - \hat{w}_1$ and σ is sufficiently small that a firm in category j chooses $\hat{\ell}_{1j}$ of low-productivity workers.

In order for the low- and high-productivity workers' intended consumption-work bundles to be unchanged, let the low-productivity workers' tax increase by $\sigma \hat{h}_1$ and the high-productivity workers' tax be unchanged. Then, the low- and high-productivity workers' incentive-compatibility constraints (3.16) and (3.17) will be slack, and since the minimum wage is the same in all firms, the low-productivity workers' incentive-compatibility constraint

(3.19) is satisfied. Also, the firms' minimum wage constraint (3.18) will be slack and the resource constraint (3.20) is satisfied. As a result, we have an equilibrium with a graduated minimum wage in which both the incentive-compatibility constraints that one type of worker does not want to mimic the other type and the minimum wage constraint are slack. In this equilibrium, the utilities of all the workers are the same as in a second-best optimum with nonlinear income taxes and redistribution toward the low-productivity workers.

Step 2: Recall that the low-productivity workers' marginal rate of substitution between consumption and leisure is less than their marginal product in a second-best equilibrium with nonlinear income taxes and redistribution toward the low-productivity workers. Therefore, letting the total demand for high-productivity workers be kept unchanged by adjusting their wage, for sufficiently small positive μ_j's there is a graduated minimum wage with $m(\hat{\ell}_{1j}+\mu_j) > \hat{w}_1+\sigma$ such that firms in category j prefer to employ $\hat{\ell}_{1j}+\mu_j$ of low-productivity workers.

Let the income tax be such that the low-productivity workers in both firm categories get to consume all their additional production. The low-productivity workers will then be better off with their new consumption-work bundles. Furthermore, let the income tax provide the high-productivity workers with the same consumption-work bundle as before. Accordingly, there is an equilibrium with a graduated minimum wage in which the low- and high-productivity workers' incentive-compatibility constraints (3.16) and (3.17) as well as the firms' minimum wage constraint (3.18) are slack, and the constraints (3.19) and (3.20) hold. In the equilibrium, all the low-productivity workers obtain a higher utility and the high-productivity workers obtain the same utility as in a second-best equilibrium without a graduated minimum wage.

Step 3: Keeping the low-productivity workers' intended minimum wage and working hours in the different categories of firms as well as the high-productivity workers' wage and working hours the same as in step two, it is feasible to transfer some of the low-productivity

workers' consumption to the high-productivity workers without violating any of the constraints (3.16)-(3.20). Consequently, there exists a graduated minimum wage which can provide a strict Pareto improvement. □

The Pareto-improving graduated minimum wage distinguishes between different categories of firms similarly to how a nonlinear income tax distinguishes between different types of workers. Consequently, the minimum wage is generally not the same for all workers, and the graduated minimum wage therefore benefits some of the low-productivity workers more than others.[11] If the graduated minimum wage in the relevant range is a hyperbolic function of the low-productivity workers' working hours, then the marginal labor cost of low-productivity workers would be the same in both firm categories while the minimum wage would be lower in the firm category with higher employment of these workers.

3.6 Conclusion

This chapter has shown that a graduated minimum wage – in contrast to a constant minimum wage – can provide a strict Pareto improvement even if the economy is initially in a second-best optimum with a nonlinear income tax that redistributes toward the low-productivity workers. The explanation is that a graduated minimum wage allows the less productive workers to benefit from second-degree price discrimination which increases their wage and income simultaneously. This reduces the attractiveness for the more productive workers to mimic the less productive workers' income and therefore loosens the formers' incentive-compatibility constraint. As a result, a minimum wage policy can be justified without appealing to a particular social welfare function.

Our model is a first step in studying the merits of a graduated minimum wage and further research is needed to generalize the framework. It would be of interest to investigate

[11] This is consistent with the fact that in many countries the minimum wage depends on firm category. Thus, the minimum wage varies with firm size in Colombia, Honduras, and several U.S. states. Similarly, it varies with the sector or occupation in Argentina, Greece, and Ireland, and with the region in Brazil, Germany, and Mexico.

how, for different social welfare functions, a graduated minimum wage interacts with the optimal income tax in affecting the consumption and working hours of the different types of workers in a social optimum. Likewise, it would be worthwhile to determine how a graduated minimum wage in a social optimum depends on the production technologies and the supply of the different types of workers.

Appendix A

Appendix to Chapter 1

A.1 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 $\varphi_{\tilde{v}x} < \varphi_v$, then $\varphi_x = \varphi_{\tilde{v}x}$.

(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

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

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

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

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 $\tilde{W}_t$

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

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, $\delta_b \tilde{W}_t(a,0,0) = N g_w(a)$ 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

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

A.2 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 COMPU-

STAT 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.

A.2.1 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.[1] 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,

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

[1]The 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.

Because I have data on the number of firms in each size bin, the proportion of firms smaller than s_v or lb_v, $\Pr(size < lb_v)$, 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.[2] 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 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.[3]

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

[3] The 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.

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 1.2 yields,

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

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 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,

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

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.[4] 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

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

However, unlike in the calibration of the other technologies, this equation has two unknowns, s_{11} 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).[5] This implies that

$$\frac{\alpha_{11}}{w_{\ell}+\lambda_{11}w_h} = 1.04\frac{\alpha_{10}}{w_{\ell}+\lambda_{10}w_h}. \tag{A.2}$$

[4] Approximately 3.9% of firms employ 50 or more workers.

[5] This 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.

The solution to equations A.1 and A.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, $f_x \in \left(\underline{f_x}, \underline{f_x} + x\right)$ and $f_{v_x} \in \left(\underline{f_{v_x}}, \underline{f_{v_x}} + x\right)$, 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.

A.2.2 Calibration Results

The parameters, other than the firm-level technological parameters, used in the baseline simulation are summarized in Table A.1.

Table A.1: 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

A.3 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.1, 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.1.
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.1 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 $\max\left\{\left(w_{\ell t} - \tilde{w}_{\ell t}\right)/w_{\ell t}, \left(w_{ht} - \tilde{w}_{ht}\right)/w_{ht}\right\} < 10^{-7}$.

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.

Appendix B

Appendix to Chapter 2

B.1 Calibration Appendix

B.1.1 Worker Distribution

Data on the worker distribution, i.e. number of employed and unemployed workers, is divided into three age bins, 15-24, 25-54 and 55 and up, for each sector. Since, I consider workers only from the age of 20 and up, I assume that all the workers in the age group 15-24 are 20 and above. Since almost everybody in Canada completes high school, this is not a big stretch.

Productivity Calibration

For the productivity calibration, I need to compute the total number of employed workers in each sector at each age and experience level, $H(a, e, s)$. If there are n employed workers in a bin spanning a years, then I assume that each age has n/a employed workers in it. In addition, I assume that a worker of age a has $a - 20$ years of experience. This essentially assumes that they worked their entire life, which is a decent approximation when the unemployment rate is not too high.

Initial Distribution

Translating the empirical worker distribution into the language of the model requires several more steps for the initial worker distribution. First, the initial worker distribution needs to account also for jobless workers. Second, the worker distribution must be manipulated slightly to ensure that there is no population growth. Finally, employed workers must be divided among the various job offers. This final step is described in the body of the paper, but it is unnecessary in order to calibrate productivity because the job offers are all equally productive and only differ in the disutility imposed on the worker.

Since the data does not include workers who are not in the labor force, I divide the number of workers in the 55 and up age bin by the labor force retention rate for old workers computed above for each sector. The added workers are included as jobless workers. This ensures that I include all workers, even those that have dropped out of the labor force. I divide workers between specific ages within age bins and between experience levels in the same way as I do in the distribution of employed workers for the productivity distribution. The only difference is that now I do so for both employed and jobless workers.

Since I abstract from population growth, each age must have the same number of workers in it. Denote by x the total number of workers, employed and jobless, divided by the 50. Since workers' life spans 50 years, if there are x workers at each age, then the population will be unity and there will be no population growth. The number of employed workers is then

$$H(a,e,s) = \frac{x * employed(a,e,s)}{total_population(a)},$$

where $total_population(a)$ is the total number of workers, employed and jobless, at age a and $employed(a,h,s)$ is the number of workers at age a with experience e in sector s who

are employed. Similarly, the number of jobless workers is taken to be

$$U(a,e,s) = \frac{x * unemployed(a,e,s)}{total_population(a)}.$$

This ensures that the population is unity and each age has equal population so that there is no population growth.

B.2 Numerical Implementation

B.2.1 Worker Values and Policy

In this subsection, I describe how to compute worker values and optimal policy. Consider a worker who has already entered a given sector, and denote the value of such a worker in period t by $V_t(a,e,j)$. If the worker is currently jobless, the value for such a worker of searching, i.e. drawing a job opportunity, is

$$V_t^d(a,e,\infty) = -c + \beta \int_{-\infty}^{\infty} V_{t+1}(a+1,e,j)\, dF(j),$$

where $F(j)$ is cumulative distribution function associated with job opportunities in the sector under consideration and the superscript d indicates that this is the value of drawing a job opportunity. The same worker's value of not searching is

$$V_t^o(a,e,\infty) = \beta V_{t+1}(a+1,e,\infty).$$

where, the superscript o indicates that the worker is out of the labor force in this period. Because a jobless worker chooses between search and retirement, the maximum of the previous two equations yield the actual value of being jobless in period t for a worker of age a

and with experience e. Thus,

$$V_t(a, e, \infty) = \max\left\{V_t^d(a, e, \infty), V_t^o(a, e, \infty)\right\}.$$

Next consider a worker with job opportunity j in period t. If the worker chooses to reject the job opportunity and instead enter the pool of jobless workers in period t, her value is the same as that of a jobless worker, $V_t(a, e, \infty)$. If instead the worker accepts the job opportunity her value is

$$V_t^w(a, e, j) = \frac{p_t A_t h(a, e)}{P_t} - c - j + \beta \tilde{V}_{t+1}^w(a, e, j),$$

where the superscript w denotes working and

$$\tilde{V}_{t+1}^w(a, e, j) = \left[(1-\lambda) V_{t+1}(a+1, e+1, j) + \lambda V_{t+1}(a+1, e+1, \infty)\right],$$

denotes the expected value of the job in the subsequent period accounting for the fact that the worker will lose her job with probability λ. Since a worker with a job opportunity must decide whether to accept or reject, the value of job opportunity j is

$$\max\left\{V_t^w(a, e, j), V_t(a, e, \infty)\right\}.$$

The final step in computing worker values is to note that $V(A+1, e, j) = 0$ for all e and j because the worker no longer exists at age $A+1$. Thus, for given paths of sectoral prices and productivities, starting for $a = A$ it is straightforward to solve backwards for worker values and policies in every period of a worker's life.

B.2.2 Worker Distribution

Suppose period t begins with the worker distribution is $W_t'(a, e, j, s)$. The prime on the distribution indicates that this is the distribution of workers before they have rejected or accepted the job opportunities that they have outstanding from the previous period. In this subsection, I show how to compute W_{t+1}' given the worker decision rules in each sector, $j_{ts}(a, e)$ and $a_{ts}(e)$, as well as the values of entering each sector, V_s. I begin by computing the distribution of newborn workers. The probability of a worker entering sector s is

$$W_t'(1, 0, \infty, s) \int_{-\infty}^{\infty} \prod_{s' \neq s} \left[1 - F_{s'}(V_{s'} - C)\right] f_s(V - C),$$

where $F_s(\cdot)$ is the CDF of the sector-entry cost draws and $f_s(\cdot)$ is the corresponding PDF. Since all entrants have no experience and are jobless, $W_t(1, e, j, s) = 0$ for all $e > 0$ and $j < \infty$. The rest of the distribution, for $a > 1$, at the beginning of period t is given by the initial W_t.

Next, all the workers with rejected job opportunities much be reintroduced into the jobless worker pool.

$$W_t(a, e, \infty, s) = \int_{j_{ts}(a,e)}^{\infty} W_t'(a, e, j, s) + W_t'(a, e, \infty, s) \qquad \forall a, e, s.$$

For $j < j_s(a, e)$, i.e. for accepted job opportunities, $W_t(a, e, j, s) = W_t'(a, e, j, s)$. Since newborn entry and job rejection can be thought of as occurring between the periods, in this case between $t-1$ and t, the distribution W_t is the distribution at the beginning of period t.

The mass of workers at job j in the subsequent period is equal to those who worked at j and retained their job plus those who drew a job offer j in period t. As a result, for $j < \infty$,

$$W_{t+1}'(a, e, j, s) = (1 - \lambda) W_t(a - 1, e - 1, j, s) + W_t(a - 1, e, j, s) k(j; a, e, s),$$

where $k(j; a, e, s)$ is the probability of drawing job j for a worker characterized by (a, e, ∞, s). This probability is endogenous because a jobless worker may decide not to search. In particular,

$$g_t(j; a, e, s) = \begin{cases} 0 & if\ a > a_{ts}(e) \\ g_s(j) & otherwise \end{cases},$$

where $g_s(j)$ is the PDF of the job-opportunity distribution in sector s.

Finally, the mass of workers who will enter the subsequent period with no job opportunity includes jobless workers who chose not to search and workers who worked but lost their job. Thus, for $j = \infty$,

$$W'_{t+1}(a, e, \infty, s) = \begin{cases} \lambda \int_{-\infty}^{j_{ts}(a,e)} W_t(a-1, e-1, j, s) + W_t(a-1, e, \infty, s) & if\ a > a_{ts}(e) \\ \lambda \int_{-\infty}^{j_{ts}(a,e)} W_t(a-1, e-1, j, s) & otherwise \end{cases}.$$

B.2.3 Simulation

Discretization

For the simulation, I discretize the job offer distribution. Since the job offer distribution is normally distributed, I partition the interval of job offers $(-\infty, \infty)$ into 100 intervals. Each interval is associated with a single job offer. The probability of drawing the job offer associated with the interval (x_1, x_2) is equal to the probability of drawing a job offer in this interval in the original normal distribution. In addition, the disutility associated with that job is equal to the expected value of the disutility draw from the original normal distribution conditional on a drawing a job offer in that interval.

For computational economy, I allow for only 10 experience levels. These experience levels correspond to five-year intervals of experience, i.e. the first corresponds to 0 to 4

years experience, the second to 5 to 9 years of experience, etc. The experience premium is equal to the empirical experience premium at the midway point of the each interval, i.e. the experience premium in the first experience level is equal to the empirical experience premium at 2.5 years of experience, etc. Finally, if a worker is jobless then she remains at her current experience level in the subsequent period. However, if she is employed she moves to the next experience level with probability 1/(52*5) and remains at her current level with the complementary probability. Since there are 52 periods in a year, this ensures that workers must work on average 5 years to move to the next experience level.

Algorithm

In this Appendix, I provide a brief description of the algorithm I use for computing an equilibrium path given the initial distribution of workers to a terminal steady state. It is necessary to simulate until a terminal steady state because in order to compute worker policies, I must know the prices they face for their entire lives. Without an endpoint this computation is impossible. This point will become clearer when the algorithms details are spelled out.

To solve for the terminal steady state, I solve for the price of the nontraded-sector good, such that the worker distribution is constant given world prices of the traded goods. This is the steady state price of services given the prices in the traded sectors. 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. If the algorithm does not converge, then a larger T is chosen.

1. Guess a path of service-sector prices $\{p_{tS}\}_{t=1}^{T}$.
2. Compute optimal worker policy given $\{p_{tS}\}_{t=1}^{T}$ by solving backward for worker values as described in Appendix B.2.1.
3. Using worker policy computed in 2 and the initial worker distribution, solve forward for the path of worker distributions, $\{W_t(a,e,j,s)\}_{t=1}^{T}$, as described in Appendix B.2.2.

4. For each period t, compute supply of and demand for services, as described in Section 1.2. If supply equals demand for all periods, then the algorithm has converged. More specifically, if $|(supply - demand)/supply| < 10^{-6}$, then the model is considered to have converged
5. If the algorithm has not yet converged then adjust the guess for the path of service prices and return to Step 2 and iterate until convergence. The guess is adjusted by finding, for each period t, the price of services $\tilde{p}_{tS}$ such that demand for services equals the supply found in 4. The demand is computed for the income in the economy implied by $W_t(a, e, j, s)$ and p_{tS}. The new guess is then equal to $\alpha p_{tS} + (1 - \alpha)\tilde{p}_{tS}$.

The parameter α must be chosen to ensure a stable convergence. Larger α's tend to speed up the convergence, but also tend to make the convergence less stable.

Appendix C

Appendix to Chapter 3

The relevant Lagrangian is[1]

$$\mathcal{L} = U(c_1, h_1) + \lambda_1 U(c_2, h_2) + \lambda_2 \left[f(\ell_1, \ell_2) - n_1 c_1 - n_2 c_2 \right] + \lambda_3 \left[U(c_2, h_2) - U\left(c_1, \frac{w_1 h_1}{w_2} \right) \right],$$

where $\lambda_1,\ \lambda_2,\ \lambda_3 > 0$. The first-order conditions for a second-best optimum with nonlinear taxes include

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial c_1} &= \frac{\partial U(c_1, h_1)}{\partial c_1} - \lambda_2 n_1 - \lambda_3 \frac{\partial U\left(c_1, w_1 h_1 / w_2\right)}{\partial c_1} = 0, \\
\frac{\partial \mathcal{L}}{\partial h_1} &= \frac{\partial U(c_1, h_1)}{\partial h_1} + \lambda_2 n_1 f_{\ell_1} - \lambda_3 \frac{\partial U\left(c_1, w_1 h_1 / w_2\right)}{\partial h_2} \left[\frac{w_1}{w_2} + h_1 \frac{\partial \left(w_1 / w_2\right)}{\partial h_1} \right] = 0.
\end{aligned}$$

Accordingly,

$$\begin{aligned}
\frac{\partial U(c_1, h_1)}{\partial c_1} &= \lambda_2 n_1 + \lambda_3 \frac{\partial U\left(c_1, w_1 h_1 / w_2\right)}{\partial c_1}, \\
\frac{\partial U(c_1, h_1)}{\partial h_1} &= -\lambda_2 n_1 f_{\ell_1} + \lambda_3 \frac{\partial U\left(c_1, w_1 h_1 / w_2\right)}{\partial h_2} \left[\frac{w_1}{w_2} + h_1 \frac{\partial \left(w_1 / w_2\right)}{\partial h_1} \right],
\end{aligned}$$

[1] The proof is based on Stiglitz (1982).

and hence

$$-\frac{\partial U(c_1,h_1)/\partial h_1}{\partial U(c_1,h_1)/\partial c_1}=\frac{\lambda_2 n_1 f_{\ell_1}-\lambda_3\left[\partial U\left(c_1,w_1h_1/w_2\right)/\partial h_1\right]\left[w_1/w_2+h_1\partial\left(w_1/w_2\right)/\partial h_1\right]}{\lambda_2 n_1+\lambda_3\partial U\left(c_1,w_1h_1/w_2\right)/\partial c_1}. \tag{A1}$$

Using that $w_1=f_{\ell_1}$ and $w_2=f_{\ell_2}$ we obtain that

$$\frac{\partial\left(w_1/w_2\right)}{\partial h_1}=\frac{\partial\left(f_{\ell_1}/f_{\ell_2}\right)}{\partial h_1}=\frac{\left(f_{\ell_1\ell_1}f_{\ell_2}-f_{\ell_1}f_{\ell_2\ell_1}\right)n_1}{f_{\ell_2}^2}<0.$$

Hence, (A1) implies that

$$-\frac{\partial U(c_1,h_1)/\partial h_1}{\partial U(c_1,h_1)/\partial c_1}<\frac{\lambda_2 n_1 f_{\ell_1}-\lambda_3\left[\partial U\left(c_1,w_1h_1/w_2\right)/\partial h_1\right]f_{\ell_1}/f_{\ell_2}}{\lambda_2 n_1+\lambda_3\partial U\left(c_1,w_1h_1/w_2\right)/\partial c_1},$$

which may be written as

$$\frac{\theta}{f_{\ell_1}}<\frac{1+\psi\phi/f_{\ell_2}}{1+\psi}, \tag{A2}$$

where

$$\begin{aligned}\theta &\equiv -\frac{\partial U(c_1,h_1)/\partial h_1}{\partial U(c_1,h_1)/\partial c_1},\\ \phi &\equiv \frac{\partial U\left(c_1,w_1h_1/w_2\right)/\partial h_1}{\partial U\left(c_1,w_1h_1/w_2\right)/\partial c_1},\\ \psi &\equiv \frac{\partial U\left(c_1,w_1h_1/w_2\right)}{\partial c_1}\frac{\lambda_3}{\lambda_2 n_1}.\end{aligned}$$

Since the agent monotonicity property implies that $\theta/f_{\ell_1}>\phi/f_{\ell_2}$, it follows from (A2) that $\theta<f_{\ell_1}$. That is, the low-productivity workers' marginal rate of substitution between consumption and leisure is less than their marginal product.

Bibliography

[1] P. Aghion, Y. Algan, P. Cahuc. Civil society and the state: The interplay between cooperation and minimum wage regulation. *Journal of the European Economic Association*, 9(1):3–42, 2011.

[2] S. P. Allen. Taxes, redistribution, and the minimum wage: a theoretical analysis. *The Quarterly Journal of Economics*, strony 477–490, 1987.

[3] J. E. Anderson, E. van Wincoop. Gravity with gravitas: A solution to the border puzzle. *The American Economic Review*, 93(1):170–192, 2003.

[4] J. E. Anderson, E. van Wincoop. Trade costs. *Journal of Economic Literature*, 42:691–751, 2004.

[5] E. Artuç. Intergenerational effects of trade liberalization. 2009.

[6] E. Artuç, S. Chaudhuri, J. McLaren. Trade shocks and labor adjustment: A structural empirical approach. *The American Economic Review*, 100(3):1008–1045, 2010.

[7] O. Ashenfelter, R. S. Smith. Compliance with the minimum wage law. *The Journal of Political Economy*, strony 333–350, 1979.

[8] A. Atkeson, A. Burstein. Innovation, firm dynamics, and international trade. Raport instytutowy, National Bureau of Economic Research, 2007.

[9] A. K. Basu, N. H. Chau, R. Kanbur. Turning a blind eye: Costly enforcement, credible commitment and minimum wage laws*. *the economic journal*, 120(543):244–269, 2010.

[10] M. Belot, J. Boone, J. Van Ours. Welfare-improving employment protection. *Economica*, 74(295):381–396, 2007.

[11] A. B. Bernard, J. Eaton, J. B. Jenson, S. Kortum. Plants and productivity in international trade. Raport instytutowy, National Bureau of Economic Research, 2000.

[12] A. B. Bernard, J. B. Jensen. Exporters, skill upgrading, and the wage gap. *Journal of International Economics*, 42(1):3–31, 1997.

[13] A. B. Bernard, J. B. Jensen, S. J. Redding, P. K. Schott. Firms in international trade. *The Journal of Economic Perspectives*, 21(3):105–130, 2007.

[14] V. Bhaskar, A. Manning, T. To. Oligopsony and monopsonistic competition in labor markets. *The Journal of Economic Perspectives*, 16(2):155–174, 2002.

[15] V. Bhaskar, T. To. Minimum wages for ronald mcdonald monopsonies: A theory of monopsonistic competition. *The Economic Journal*, 109(455):190–203, 1999.

[16] T. Blumkin, E. Sadka. Income taxation and wage policy: an application to minimum wage. *International Tax and Public Finance*, 12(6):713–722, 2005.

[17] R. Boadway, K. Cuff. A minimum wage can be welfare-improving and employment-enhancing. *European Economic Review*, 45(3):553–576, 2001.

[18] T. Boeri. Setting the minimum wage. *Labour Economics*, 19(3):281–290, 2012.

[19] T. Boeri, M. C. Burda. Preferences for collective versus individualised wage setting*. *The Economic Journal*, 119(540):1440–1463, 2009.

[20] W. Brown. The process of fixing the british national minimum wage, 1997–2007. *British Journal of Industrial Relations*, 47(2):429–443, 2009.

[21] A. Burstein, J. Vogel. Globalization, technology, and the skill premium: A quantitative analysis. Raport instytutowy, National Bureau of Economic Research, 2010.

[22] P. Bustos. The impact of trade on technology and skill upgrading evidence from argentina. *Working Papers (Universitat Pompeu Fabra. Departamento de Economía y Empresa)*, (1189):1, 2007.

[23] P. Bustos. Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *The American Economic Review*, 101(1):304–340, 2011.

[24] P. Cahuc, G. Laroque, i in. Optimal taxation and monopsonistic labor market: does monopsony justify the minimum wage? 2009.

[25] P. Cahuc, A. Saint-Martin, A. Zylberberg. The consequences of the minimum wage when other wages are bargained over. *European Economic Review*, 45(2):337–352, 2001.

[26] R. M. Campos-Vázquez. Why did wage inequality decrease in mexico after nafta? Raport instytutowy, El Colegio de México, Centro de Estudios Económicos, 2010.

[27] T. Chaney. Distorted gravity: The intensive and extensive margins of international trade. *The American Economic Review*, 98(4):1707–1721, 2008.

[28] D. Checchi, C. García-Peñalosa. Labour market institutions and income inequality. *Economic Policy*, 23(56):601–649, 2008.

[29] K. Cosar. Adjusting to trade liberalization: Reallocation and labor market policies. *V manuscript*, 2010.

[30] J. Costantini, M. Melitz. The dynamics of firm-level adjustment to trade liberalization. *The Organization of Firms in a Global Economy*, strony 107–141, 2007.

[31] L. Danziger. The elasticity of labor demand and the minimum wage. *Journal of Population Economics*, 22(3):757–772, 2009.

[32] L. Danziger. Endogenous monopsony and the perverse effect of the minimum wage in small firms. *Labour Economics*, 17(1):224–229, 2010.

[33] C. Davidson, L. Martin, S. Matusz. Trade and search generated unemployment. *Journal of International Economics*, 48(2):271–299, 1999.

[34] C. Davidson, S. J. Matusz. Globalization and labour-market adjustment: how fast and at what cost? *Oxford Review of Economic Policy*, 16(3):42–56, 2000.

[35] D. R. Davis, J. Harrigan. Good jobs, bad jobs, and trade liberalization. 2007.

[36] G. De Fraja. Minimum wage legislation, productivity and employment. *Economica*, 66(264):473–488, 1999.

[37] H. Egger, U. Kreickemeier. Firm heterogeneity and the labor market effects of trade liberalization*. *International Economic Review*, 50(1):187–216, 2009.

[38] P. Fajgelbaum, G. M. Grossman, E. Helpman. Income distribution, product quality, and international trade. *Journal of Political Economy*, 119(4):721–765, 2011.

[39] R. Falvey, D. Greenaway, J. Silva. Trade liberalisation and human capital adjustment. *Journal of International Economics*, 81(2):230–239, 2010.

[40] R. Findlay, H. Kierzkowski. International trade and human capital: A simple general equilibrium model. *The Journal of Political Economy*, strony 957–978, 1983.

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

[42] R. B. Freeman. A cobweb model of the supply and starting salary of new engineers. *Industrial and Labor Relations Review*, 29(2):236–248, 1976.

[43] P. K. Goldberg, N. Pavcnik. Distributional effects of globalization in developing countries. *Journal of Economic Literature*, 45:39–82, 2007.

[44] R. Guesnerie, K. Roberts. Minimum wage legislation as a second best policy. *European Economic Review*, 31(1):490–498, 1987.

[45] J. Harrigan, A. Reshef. Skill biased heterogeneous firms, trade liberalization, and the skill premium. Raport instytutowy, National Bureau of Economic Research, 2011.

[46] R. G. Harris, P. E. Robertson. Trade, wages and skill accumulation in the emerging giants. *Journal of International Economics*, 89(2):407–421, 2013.

[47] A. Harrison, J. McLaren, M. McMillan. Recent perspectives on trade and inequality. *Annual Review of Economics*, 3(1):261–289, 2012.

[48] J. J. Heckman, L. Lochner, C. Taber. 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, 1998.

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

[50] M. Hungerbühler, E. Lehmann. On the optimality of a minimum wage: New insights from optimal tax theory. *Journal of Public Economics*, 93(3):464–481, 2009.

[51] L. F. Katz, K. M. Murphy. Changes in relative wages, 1963–1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78, 1992.

[52] D.-I. Kim, R. H. Topel. Labor markets and economic growth: Lessons from korea's industrialization, 1970-1990. *Differences and Changes in Wage Structures*, strony 227–264. University of Chicago Press, 1995.

[53] N. Kwark, C. Rhee. Educational wage differentials in korea. 1993.

[54] D. Lee, E. Saez. Optimal minimum wage policy in competitive labor markets. *Journal of Public Economics*, 96(9):739–749, 2012.

[55] A. Lileeva, D. Trefler. Improved access to foreign markets raises plant-level productivity... for some plants. *The Quarterly Journal of Economics*, 125(3):1051–1099, 2010.

[56] A. Manning. *Monopsony in motion: Imperfect competition in labor markets*. Princeton University Press, 2003.

[57] N. Marceau, R. Boadway. Minimum wage legislation and unemployment insurance as instruments for redistribution. *The Scandinavian Journal of Economics*, strony 67–81, 1994.

[58] K. Matsuyama. A simple model of sectoral adjustment. *The Review of Economic Studies*, 59(2):375–387, 1992.

[59] K. Matsuyama. Beyond icebergs: Towards a theory of biased globalization. *The Review of Economic Studies*, 74(1):237–253, 2007.

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

[61] G. G. Porto. Using survey data to assess the distributional effects of trade policy. *Journal of International Economics*, 70(1):140–160, 2006.

[62] J. B. Rebitzer, L. J. Taylor. The consequences of minimum wage laws some new theoretical ideas. *Journal of Public Economics*, 56(2):245–255, 1995.

[63] R. Robertson. Trade and wages: Two puzzles from mexico. *The World Economy*, 30(9):1378–1398, 2007.

[64] B. Salanie. *The economics of taxation*. MIT press, 2011.

[65] R. S. Sobel. Theory and evidence on the political economy of the minimum wage. *Journal of Political Economy*, 107(4):761–785, 1999.

[66] M. B. Stewart, J. K. Swaffield. The other margin: Do minimum wages cause working hours adjustments for low-wage workers? *Economica*, 75(297):148–167, 2008.

[67] J. E. Stiglitz. Self-selection and pareto efficient taxation. *Journal of Public Economics*, 17(2):213–240, 1982.

[68] E. Strobl, F. Walsh. The ambiguous effect of minimum wages on hours. *Labour Economics*, 18(2):218–228, 2011.

[69] R. H. Topel. Factor proportions and relative wages: The supply-side determinants of wage inequality. *The Journal of Economic Perspectives*, 11(2):55–74, 1997.

[70] J. R. Tybout. Plant-and firm-level evidence on "new" trade theories. *Integration & Technology*, strona 67.

[71] E. A. Verhoogen. Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. *The Quarterly Journal of Economics*, 123(2):489–530, 2008.

[72] D. Weil. Public enforcement/private monitoring: Evaluating a new approach to regulating the minimum wage. *Industrial and Labor Relations Review*, strony 238–257, 2005.

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