Essays in International Trade and Development

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Michael Maio

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

TIMOTHY KEHOE AND FABRIZIO PERRI

JULY, 2014

Acknowledgements

I owe a debt of gratitude to many people for the role they played in helping me obtain a Ph.D. Timothy Kehoe and Fabrizio Perri have provided invaluable guidance as advisers. Jim Schmitz supported me as his research assistant at the Federal Reserve Bank of Minneapolis and taught me never to accept data at face value, and to always try to understand at a deep level how it is created. Enoch Hill and David Perez-Reyna, my co-authors on the second chapter, have been great colleagues and friends, and our essay would not have been possible without them. All of the participants in the University of Minnesota Trade and Development Workshop from 2010-2014 patiently listened to me present my research about 9,000 times and provided considered feedback. The comments and suggestions provided by Thomas Holmes, Cristina Arellano, Arilton Teixeira, and Jose Asturias also deserve to be singled out.

Dedication

For my parents, Douglas and Mary Maio.

Abstract

In the fields of international trade and economic development, the behaviors of individual firms are often at the center of attention. What unites the essays in this thesis is their focus on studying the aggregate effects of individual firm decisions. Chapter 2 demonstrates that pressure from foreign competition can reduce managerial slack in domestic firms, and that this was a substantial source of productivity gains in Chile following the country's unilateral trade liberalization in the 1970s. Chapter 3 demonstrates theoretically how firms' choices of organizational structure lead to the high skill premium seen among large firms. Finally, in Chapter 4, I show that in an environment where firm market power differs across industries, firm technological upgrading decisions amplify market distortions and increase allocative inefficiency.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The essays in this thesis fall broadly within the fields of international trade and economic development. In both fields, the behaviors of individual firms are often at the center of attention, and what unites the essays in this thesis is their focus on studying the aggregate effects of individual firm decisions.

Firms in import-competing industries often experience substantial productivity growth following unilateral trade liberalization episodes. In Chapter 2, I demonstrate that a firm's ownership structure plays an important role in understanding the relationship between foreign competition and firm productivity growth. The central idea is that foreign competition leads to productivity gains in firms with separate ownership and management by reducing moral hazard. That is, foreign competition reduces managerial slack. I develop and calibrate a model to study this mechanism, and the mitigation of principal-agent frictions emerges as a quantitatively important source of productivity gains among firms in import-competing industries. When applied to Chile's unilateral tariff reduction in the 1970s, the model accounts for 69 percent of the difference in average productivity growth between plants in import-competing industries and those in non-traded goods industries. The model also accounts for the fact that, in Chile's import-competing industries, plants with separate ownership and management had much higher productivity growth than owner-managed plants.

Chapter 3 is co-written with Enoch Hill and David Perez-Reyna. In it, we propose an original model of firm hierarchy which suggests that a firm's organizational structure is important for understanding its wage structure. In our model, more productive firms choose to employ more levels of management, which requires a higher average level of skill in workers and consequently a higher average skill premium. This is consistent both with our findings from a Chilean firm-level data set and also with the positive relationship between firm size and skill premium documented in the literature. Additionally, our model predicts that skill premium is increasing in the ratio of workers to managers, a fact we also observe in the Chilean data.

In Chapter 4, I step outside the trade context of Chapter 2 and study theoretically the relationship between product market competition and innovation in a general equilibrium environment in which different industries have different levels of competitiveness, as measured by the number of producers. In particular, I develop a model to study how firm technological upgrading decisions affect allocative efficiency in a multi-industry environment where market concentration differs across industries. I conduct numerical simulations for a range of empirically relevant parameters, and I find that technological upgrading increases the allocative inefficiency that results from differences in market power across industries. Firms in industries with a low number of producers and high market concentration do not upgrade their productive technology, while firms in more competitive industries do. Resources then shift away from industries that are already producing too few goods relative to the socially optimal level. Allocative inefficiency increases. Once that firm upgrading decisions are taken into account, the efficiency losses from differences in market concentration become even larger.

Chapter 2

Foreign Competition and Firm Productivity: A Principal-Agent Approach

2.1 Introduction

Many empirical studies show that firms in import-competing industries often experience substantial productivity growth following a unilateral trade liberalization. When tariffs fall, domestic firms face increased competitive pressure from abroad, and many respond by raising their productivities.[1] In this essay, I demonstrate both theoretically and empirically that a firm's ownership structure plays an important role in understanding the relationship between firm productivity growth and import pressure. Using a plant-level manufacturing dataset from Chile, I document that plants with separate ownership and management (e.g., corporations) in import-competing industries had high average productivity growth following the country's unilateral tariff reduction in the 1970s. Meanwhile, owner-managed plants (e.g., sole proprietorships) in import-competing industries did not. Existing theories of trade and firm-level productivity growth ignore ownership structure, but the data call for a theory that does not.

[1] The growing list of countries where this phenomenon has been documented includes Chile (Pavcnik 2002), Colombia (Fernandes 2007), India (Sivadasan 2009), and Brazil (Muendler 2004).

The central theoretical idea of this paper is that foreign competition leads to productivity gains in firms with separate ownership and management by ameliorating moral hazard. Put simply, foreign competition reduces managerial slack. I develop and calibrate a model to study this mechanism, and the mitigation of principal-agent frictions emerges as a quantitatively important source of productivity gains among firms in import-competing industries.

The model is also consistent with the high productivity growth seen in Chilean import-competing industries among plants with separate ownership and management, relative to owner-managed plants. In my model, some firms are owner-managed and some have separate ownership and management. Managers of all firms (regardless of ownership type) work harder to raise productivity when foreign competition threatens the firm's survival. The reason is that they value their jobs and do not want to lose them. A critical point is that the effect of competition on managerial effort is much stronger in the presence of moral hazard. Increased foreign competition leads to high productivity growth in firms with separate ownership and management largely by relieving principal-agent frictions. In owner-managed firms, where there is no managerial slack, there is less scope for productivity gains.

In my model, the fear of bankruptcy leads to productivity improvements. By contrast, the standard explanation for why trade liberalization leads to firm-level productivity growth involves scale economies. Prominent examples include Costantini and Melitz (2007), Lileeva and Trefler (2010), and Bustos (2011). Their models are designed to study productivity growth among exporting firms following bilateral liberalizations, but they are not appropriate for studying unilateral liberalizations. They would predict that a unilateral fall in tariffs would *decrease* rather than *increase* productivity growth among firms in import-competing industries. The logic of scale economies is as follows. As trade barriers fall bilaterally, exporting firms sell a larger quantity of goods in foreign markets, which increases the gains associated with productivity improvements and therefore leads to productivity growth. Scale economies cannot explain why firms in import-competing industries see productivity growth following unilateral liberalizations, though: their sales shrink rather than increase.

As mentioned above, there is need for a theory that not only captures productivity gains in import-competing industries, but also captures differences by ownership type.

The differences in the Chilean data are quite large. During the period of adjustment to Chile's unilateral tariff reduction in the late 1970s, plants in import-competing industries experienced productivity growth that was 8.9 percentage points higher on average than plants in non-traded goods industries. One contribution of my paper is to show that most of the difference can be attributed to high productivity growth among the plants in import-competing industries that I categorize as having principal-agent frictions. More precisely, I divide plants into two groups based on ownership type. One group, which I classify as *managerial* plants, consists of plants with separate ownership and management. These are the plants afflicted by principal-agent frictions. Corporations account for the largest share of managerial plants in the data. The other group, *entrepreneurial* plants, consists of plants in sole proprietorships or partnerships, where there is no separation between ownership and management.

Managerial plants in import-competing industries have average productivity growth 16.5 percentage points higher than plants in non-traded industries (of either ownership type). Meanwhile, entrepreneurial plants in import-competing industries have average productivity growth that is roughly the same as plants in non-traded industries. Within non-traded industries, there is no statistically significant difference in average productivity growth between ownership types. The results suggest that competitive pressure leads to productivity gains in part by alleviating principal-agent frictions in firms with separate ownership and management.

To investigate the extent to which the reduction of principal-agent frictions in firms can explain the high average productivity growth seen in import-competing industries, I nest a moral hazard problem in a multi-industry heterogeneous firm trade model. The model borrows elements from Bernard, Eaton, Jensen, and Kortum (2003) and, in addition, contains both entrepreneurial and managerial firms. I model entrepreneurial firms as sole proprietorships whose owner-managers receive all firm profits. Managers of managerial firms receive a share of profits agreed to through a contract with the firm owners.

Managers of all firms can raise productivity by a fixed amount by exerting effort. Exerting more effort increases the probability that the manager successfully raises productivity. Managers also receive a non-pecuniary benefit from running their firms. When the survival of a firm is threatened by a foreign competitor, the non-pecuniary benefit serves as an incentive for exerting effort to raise productivity and keep the firm profitable. Even though foreign competition reduces the demand for a firm's good - and therefore the pecuniary rewards from raising productivity - the threat of losing the non-pecuniary benefit can induce a manager to exert greater effort following the increase in foreign competition. In this way, the "bankruptcy effect" of competition on productivity more than offsets the standard scale effect.

Owners and managers in managerial firms face a standard moral hazard problem in the class of problems studied in Sappington (1982). The two parties, who are both risk-neutral, cannot form a contract that induces the manager to exert effort at the first-best level because of a limited liability constraint on the manager. The contract that induces the manager to exert effort at the first-best level requires the manager to receive negative payments when productivity is low, which I do not allow. Managers therefore slack and exert effort below the first-best level when ownership and management are separate.

In the model's import-competing industry, as in the data, managerial firms see higher average productivity growth than entrepreneurial firms. The reason is that managers of entrepreneurial firms, who internalize all the pecuniary gains from raising productivity, exert effort at a high level even before tariffs fall. Because the disutility of effort is convex, it is highly costly for them to increase their effort levels even further once foreign competition increases. In managerial firms, which have separate ownership and management, managers slack when faced with little competitive pressure. As a result, they can exert more effort at comparatively little cost in response to increased foreign competition.

When applied to Chile's trade reforms, the model accounts for most of the difference in average productivity growth between plants in import-competing industries and plants in non-traded industries. I calibrate the model to data from Chile and simulate the country's tariff reduction. I find that the difference in average productivity growth between the model's import-competing industry and the model's non-traded industry is 6.1 percent. In the data, the figure is 8.9 percent, so the mechanism I study can account for 69 percent of the difference in average growth.

The model also generates quantitatively important differences in productivity growth by ownership type in the import-competing industry. In the data, import-competing managerial plants have average productivity growth 16.5 percentage points higher than import-competing entrepreneurial plants. In the model, the difference is 5.6 percentage points. The lessening of principal-agent frictions in the model therefore accounts for 34 percent of the difference in average productivity growth between the two ownership types in the data. Hence the reduction of principal-agent frictions is an important contributor to the high productivity growth seen among import-competing managerial plants. In both the model and the data, differences in productivity growth between both managerial and entrepreneurial producers of non-traded products are negligible.

I also use the model to analyze the aggregate effects of foreign competition on managerial effort. The policy simulation in the model leads to a 1.6 percent increase in real income. Of the 1.6 percent increase, 0.3 percentage points are attributable to the alleviation of principal-agent frictions, so inefficiencies due to moral hazard problems within firms represent a quantifiable and economically meaningful source of aggregate loss. Removing trade barriers serves to reduce the losses from principal-agent frictions by aligning the incentives of owners and managers more closely.

The model in this paper is not the only one to focus on the effects of foreign competition on the productivities of import-competing firms. Holmes and Schmitz (2001) show that when firms face foreign competitors, they may invest resources into raising productivity rather than using resources to block the innovative activities of domestic competitors. In the trapped factor model of Bloom, Romer, Terry, and Van Reenen (2013), greater competitive pressure from low-wage countries induces firms in high-wage countries to shift resources away from producing "old goods" and toward innovation. Others studies, such as Holmes, Levine, and Schmitz (2012), examine the theoretical effect of competition on productivity in environments that do not explicitly include international trade, but their key ideas could be extended to explain why foreign competition leads import-competing firms to raise productivity. I view the above studies as being complementary to mine. My paper differs in that I am able to address the relationship between ownership structure and productivity growth that I see in the Chilean data.

The principal-agent problem at the heart of my model borrows from a well-established literature on moral hazard. As in Harris and Raviv (1979), the incentive problem between the manager and owners arises from the fact that it is costly for the manager to exert unobservable effort to raise productivity, but that owners are indifferent to the manager's level of effort, all else equal. Schmidt (1997) and Raith (2003) investigate the theoretical effect of increased product market competition on managerial effort in settings with moral hazard and show conditions under which increased competition leads to increased effort. The form of the principal-agent problem in my model is closely related to the problem considered in Schmidt (1997), although Schmidt's analysis considers a partial equilibrium environment with no international trade. In contrast to previous work, my paper presents a calibrated model with international trade and evaluates the quantitative importance of principal-agent frictions in explaining the relationship between foreign competition and productivity growth.

The rest of the essay proceeds as follows. Section 2.2 presents the model, and Section 2.3 characterizes the equilibrium. Section 2.4 provides evidence from Chile in support of the model's main mechanism. Section 2.5 describes the calibration strategy and presents the main quantitative results of the paper. Section 2.6 quantifies the role of principal-agent frictions in the model and discusses the results, and Section 2.7 concludes.

2.2 Model

Here I present a two-country, multi-industry model with heterogenous firms that can be used to study firm productivity dynamics following Chile's trade liberalization. Similar to Bernard, Eaton, Jenson, and Kortum (2003), there is a continuum of imperfectly substitutable traded products, and the producer of a particular traded product competes head-to-head with a foreign producer of the same product. Additionally, each country has a continuum of non-traded products, which are produced by firms that behave as monopolists. The main theoretical innovation is to embed a principal-agent problem within a heterogeneous-firm trade model. Firms are operated by managers who can raise the productivity of a firm by exerting effort. The firms' owners (principals) use performance incentives to induce the managers to exert effort, and as we will see, a manager will exert a high level of effort if doing so will prevent the firm from shutting down.

2.2.1 Environment

There are two countries, denoted by h and f. The labor supplies are L_h and L_f, respectively. I assume $L_f >> L_h$, so the home country is much smaller than the foreign country. I also assume that the home country is too small to affect aggregate equilibrium quantities in the foreign country, but it still faces a downward-sloping foreign demand curve for its products. In the quantitative analysis, Chile will play the role of the small home country, while the large foreign country will represent the rest of the world. In each country there are three industries, a non-traded industry and two traded industries. Within each country i and industry k there is a mass M_{ik} of firms in operation, each of which produces a differentiated product. Let the industries be indexed by $k \in \{N, I, E\}$. The letters will stand for "non-traded," "import-competing," and "export-oriented."

As will be discussed further below, the two traded industries are distinguished by the productivity distributions from which firms in the two industries draw. Capturing a notion of Ricardian comparative advantage in the model, domestic firms in the export-oriented industry draw from a productivity distribution that has a high mean compared to the distribution from which foreign firms in the same industry draw. Meanwhile, domestic firms in the import-competing industry draw from a productivity distribution with a lower mean than their foreign counterparts.

Firms in the non-traded industry also draw from an industry-specific productivity distribution. The non-traded industry should be thought of as representing manufacturing industries, such as concrete, in which transportation costs are so high that the volume of international trade is negligible. Rather than explicitly allowing transportation costs to vary by industry in the model, I simply impose that products produced in the non-traded industry cannot cross the international border.

All producers of traded products engage in head-to-head competition with a foreign producer. That is, each producer of a traded product faces competition from a single foreign foreign firm that produces a perfectly substitutable product. Consumers in each location have linear preferences over a final good, which is a CES aggregate of the differentiated products. Preferences in the home country are given by

$$c_h = \Big(M_{hN} \int_0^1 q_{hN}(j)^{\frac{\sigma-1}{\sigma}} dj + M_I \int_0^1 (q_{hhI}(j) + q_{hfI}(j))^{\frac{\sigma-1}{\sigma}} dj + M_E \int_0^1 (q_{hhE}(j) + q_{hfE}(j))^{\frac{\sigma-1}{\sigma}} dj \Big)^{\frac{\sigma}{\sigma-1}}$$

Here, $q_{hN}(j)$ is the quantity of non-traded product j consumed in the home country. The term $q_{hhk}(j)$ represents the quantity of the jth product in traded industry k that is produced in the home country and consumed at home, and $q_{hfk}(j)$ represents the quantity of the jth product in traded industry k that is produced abroad and consumed at home. Note that the masses of firms in the traded industries are not indexed by location, i. That is because each traded product is produced by exactly one firm in each location. Preferences in the foreign country are defined similarly, but with $M_{fN} >> M_{hN}$. This assumption will make the traded industries negligibly small relative to the foreign country's non-traded industry, consistent with the home country being small. I will also keep the ratio of (potential) firms to workers the same in each location:

$$\frac{M_{hN} + M_{hI} + M_{hE}}{L_h} = \frac{M_{fN} + M_{fI} + M_{fE}}{L_f}$$

To avoid redundancy, the exposition of the rest of the model will be done from the home country's perspective. Standard arguments yield the following inverse demand functions:

$$p_{hk}(j) = P_h \Big(\frac{c_h}{q_{hk}(j)}\Big)^{\frac{1}{\sigma}} \tag{2.1}$$

where the ideal price index P_h is given by

$$P_h = \Big(M_{hN} \int_0^1 p_{hN}(j)^{1-\sigma} dj + M_I \int_0^1 p_{hI}(j)^{1-\sigma} dj + M_E \int_0^1 p_{hE}(j)^{1-\sigma} dj\Big)^{\frac{1}{1-\sigma}}$$

The price $p_{hk}(j)$ is the price charged to a consumer in the home country for product j in industry k, and $q_{hk}(j)$ is the total quantity consumed of that product. (Note that $q_{hk}(j) = q_{hhk}(j) + q_{hfk}(j)$ in traded industries.) P_h is the price of the final good c_h.

2.2.2 Production

Each firm produces output with a linear labor-only technology with idiosyncratic labor productivity $z_{hk}(j)$, which is drawn from a country- and industry-specific distribution G_{hk}. Draws are i.i.d., and the draws of head-to-head competitors are uncorrelated. Additionally, a firm in the home country faces an iceberg transportation cost D_{hf} of selling goods in the foreign market. In order to study unilateral reductions in trade barriers, I allow the iceberg cost of selling to the foreign country to differ from the iceberg cost faced by foreign firms selling to the home country. That is, I will consider cases with $D_{hf} \neq D_{fh}$, where D_{fh} is the iceberg cost faced by foreign firms selling to the home country. The technology of a home country producer is summarized by

$$q_{hhk}(j) = z_{hk}(j)\ell_{hhk}(j) \tag{2.2}$$

$$q_{hfk}(j) = \frac{z_{hk}(j)\ell_{hfk}(j)}{D_{hf}} \tag{2.3}$$

The firm demands $\ell_{hhk}(j) + \ell_{hfk}(j)$ units of production labor, which are compensated with a wage w_h.

It is worth emphasizing here that a firm in the import-competing industry may export if it receives an unusually high productivity draw. Meanwhile, a firm in the export-oriented industry may find itself usurped in the home market by a foreign competitor if it receives a sufficiently low productivity draw. The two industries are distinguished by their average productivity draws, so the export-oriented industry will have a high proportion of exporters compared to the import-competing industry.

2.2.3 Managers

Firms are owned collectively by all agents. In order to operate, each home country firm must hire a single manager. Managers, who are identical, are drawn from a pool of mass $N_h \subset L_h$ and randomly matched with a firm.[2] Managers cannot switch firms but retain the option of working as production laborers instead. If they remain with their firms, they receive a share of profits specified by a contract to be discussed below.

Each manager has two roles. The first is to choose the firm's output level and price given $z_{hk}(j)$. As will become evident later, managers always choose the profit-maximizing quantities given z because their compensation increases with the firm's profits. The manager's second role - and the one that is most important in the model - is to influence the firm's productivity level by exerting effort. By exerting effort $e(z, \rho)$, the manager of a firm with an initial productivity draw of z can raise productivity to

[2] I assume that N_h is large enough that the number of firms in operation is not constrained by the supply of managers.

Δz (where $\Delta > 1$) with probability ρ. The parameter Δ is a fixed quantity given by nature. Managerial effort only affects ρ.

In what ways can a manager influence productivity? In this paper, I view the potential productivity gains embodied in Δ primarily as the one-time implementation of best business practices. Bloom, Draca, and Van Reenen (2011) find evidence that European firms raised productivity in response to increased import competition from China following the country's entry into the World Trade Organization in 2001. They show that one contributor to productivity growth was the adoption of improved business practices.[3] Schmitz (2005) highlights a related mechanism through which competition leads to productvitiy gains. He finds that competitive pressure from Brazilian iron ore producers induced American producers to abandon inefficient work rules bargained for by labor unions. While the increased competition made unions more willing to accept the changes, it also presumably made management more willing to exert effort to take on the unions. Finally, Aggarwal and Samwick (2006) find evidence that managers incur private costs when making investment decisions on behalf of firms, leading to underinvestment in technologies that raise productivity. Although the model here does not include capital or R&D goods, the main insights could be extended readily to a model that includes a managerial investment decision.

In the model, foreign firms do not have managers, and their productivities are given by their initial draws. The assumption that foreign firms cannot raise their productivities means that the model will not be able to capture any possible strategic interactions between the managers of two head-to-head competitors. While that is a potentially interesting area for future research, it is beyond the scope of this paper. Here I will assume that the firms in the rest of the world have already been subject to sufficient market discipline that they have little room for further gains by the time they compete with the home country firms.

Throughout this paper, I will let the function relating effort to productivity gains take the following form:

$$e(z, \rho) = Az^{\sigma-1}exp(b\rho) \tag{2.4}$$

[3] In Bloom and Van Reenen (2007), the authors survey firms on their management practices related to, for example, incentive provision and monitoring. They find that the ones they deem to be "better business practices" are positively correlated with firm outcomes such as profits and productivity.

where $A > 0$ is a scaling parameter and $b > 0$ governs the curvature of e with respect to ρ. Effort is increasing in the initial productivity draw z, so it is more difficult for the manager of a highly productive firm to raise z by a factor of Δ than it is for the manager of a less-productive firm to do so. The effort function is also increasing and convex in the choice of ρ. The above functional form is the same as the one used in Atkeson and Burstein (2010), except that in their model R&D labor, not managerial effort, raises productivity. As will be discussed in Section 2.5.1, setting the curvature parameter on z equal to $\sigma - 1$ makes the manager's effort decision homothetic in z for non-traded firms, so that all non-traded firms in a particular location will have the same probability ρ of having the manager's efforts be successful. Since size and productivity are perfectly correlated in the model, Gibrat's Law holds for producers of non-traded products.[4]

Finally, a manager enjoys a non-pecuniary benefit $\lambda z^{\sigma-1}$ from working as a manager at the firm, where z is the firm's initial productivity draw. The benefit can be thought of as job perquisites, or private benefits that managers receive from controlling the firm. The manager's benefit increases in the firm's initial productivity z, capturing the idea that the manager of a large, highly productive organization enjoys more benefits than the manager of a smaller, less-productive organization. I assume that firm owners cannot affect the manager's perks. The benefit $\lambda z^{\sigma-1}$ can be thought of as capturing the prestige of being the boss, political influence, or other rewards that do not directly use firm resources. The curvature parameter on z is set equal to $\sigma - 1$ for the same reason that it is set to $\sigma - 1$ in the effort function: Doing so makes the manager's choice of ρ invariant to a firm's initial productivity draw if the firm is a producer of a non-traded product.

Later I will add to the model a type of firm that is owned solely by its manager. Doing so will allow me to study the quantitative effect of ownership structure on productivity growth. Managers of the manager-owned firms have the same perks and the same opportunity to raise productivity as the managers who do not own their firms, and they

[4]The same will not be true for producers of traded products because the efforts of the managers of these firms will depend on how productive the firms are relative to their foreign competitors. However, the effort decisions of managers of traded firms will be homothetic in the following sense: A manager of a firm with productivity z facing a competitor with productivity z^* will choose the same value of ρ as the manager of a firm with productivity αz facing a foreign firm with productivity αz^*, where $\alpha > 0$.

face the same potential threats from foreign competitors. The only difference is that managers who own their own firms receive all of the firm's profits rather than just a share. The logic of the model extends readily to cover manager-owned firms, so they will not be discussed again in this section.

2.2.4 Firm profits

In the next subsection, I discuss the optimal profit-sharing contract between owners and managers, which will govern the choice of managerial effort. But for now, take the *ex-post* distributions of productivities as given. (That is, suppose the outcomes of managers' efforts are already known.) At this point it is convenient to identify a firm by its idiosyncratic productivity z and drop the index j.

Producers of non-traded products

Given productivity z, the manager of a firm that produces a non-traded product chooses output produced for the home market to maximize profits $\pi_{hN}(z)$:

$$\text{Max } \pi_{hN}(z) = p_{hN}(z)q_{hN}(z) - \frac{w_h \ell_{hN}(z)}{z}$$

$$\text{s.t. inverse demand function (1), production function (2)}$$

The producer's price is a constant markup over marginal cost:

$$p_{hN}(z) = \left(\frac{\sigma}{\sigma - 1}\right)\left(\frac{w_h}{z}\right) \tag{2.5}$$

Variable profits are then given by

$$\pi_{hN}(z) = \left(\frac{1}{\sigma}\right)\left(\frac{\sigma - 1}{\sigma}\right)^{\sigma-1} P_h^{\sigma} c_h \left(\frac{z}{w_h}\right)^{\sigma-1} \tag{2.6}$$

Producers of traded products

The two producers of a particular traded product (one domestic, one foreign) are Bertrand competitors. Since they produce the same product, the producer who supplies the product at the lowest price in a particular location captures the entire market for the product in that location. By standard arguments, the producer that can supply a location at the lowest unit cost will capture the entire market and will set a price equal to either the competitor's marginal cost (that is, it will limit-price) or the monopoly price, whichever is lower. For example, suppose a home country firm has productivity z and its foreign competitor has productivity z^*. Also suppose that $w_h/z < (D_{fh}w_f)/z^*$, so that the home firm has the lowest unit cost in its home market. The home firm captures the home market and sets the price equal to

$$p_{hk}(z, z^*) = min\Big\{\Big(\frac{\sigma}{\sigma - 1}\Big)\Big(\frac{w_h}{z}\Big), \Big(\frac{D_{fh}w_f}{z^*}\Big)\Big\} \tag{2.7}$$

Note that in the traded industries it is necessary to index the price by the productivities of both the home producer and the foreign producer. A home country producer that captures the home market and sets the monopoly price will earn profits in the home market equal to

$$\pi_{hhk}(z, z^*) = \Big(\frac{1}{\sigma}\Big)\Big(\frac{\sigma - 1}{\sigma}\Big)^{\sigma-1} P_h^{\sigma} c_h \Big(\frac{z}{w_h}\Big)^{\sigma-1} \tag{2.8}$$

And if it behaves as a limit-pricer in its home market, it will earn home-market profits equal to

$$\pi_{hhk}(z, z^*) = \Big[\phi_h^{\sigma-1} - \phi_h^{\sigma}\Big] P_h^{\sigma} c_h \Big(\frac{z}{w_h}\Big)^{\sigma-1} \tag{2.9}$$

where ϕ_h is the ratio of the home and foreign firms' unit labor costs for supplying the home market:

$$\phi_h \equiv \Big(\frac{w_h}{z}\Big)\Big/\Big(\frac{D_{fh}w_f}{z^*}\Big).$$

If the home firm is also the lowest-cost producer in the foreign market and can set the monopoly price abroad, its profits from selling abroad will be

$$\pi_{hfk}(z, z^*) = \Big(\frac{1}{\sigma}\Big)\Big(\frac{\sigma - 1}{\sigma}\Big)^{\sigma-1} P_f^{\sigma} c_f \Big(\frac{z}{D_{hf}w_h}\Big)^{\sigma-1} \tag{2.10}$$

And finally, if the home firm sells abroad and sets the limit price in the foreign market, its profits from selling abroad will be

$$\pi_{hfk}(z, z^*) = \Big[\phi_f^{\sigma-1} - \phi_f^{\sigma}\Big] P_f^{\sigma} c_f \Big(\frac{z}{D_{hf}w_h}\Big)^{\sigma-1} \tag{2.11}$$

where ϕ_f is the ratio of the home and foreign firms' unit labor costs for supplying the foreign market:

$$\phi_f \equiv \left(\frac{D_{hf}w_h}{z}\right)\Big/\left(\frac{w_f}{z^*}\right).$$

For a home country firm in traded industry k with productivity z, total profits are given by the sum of profits earned in each location: $\pi_{hk}(z, z^*) = \pi_{hhk}(z, z^*)+\pi_{hfk}(z, z^*)$. Not all producers of traded products can produce profitably. The home firm will not be the low-cost producer in either location and will shut down if

$$\frac{z}{w_h} < \frac{z^*}{D_{fh}w_f}$$

The outcome of head-to-head Bertrand competition between a home firm and its foreign rival can be summarized in the following definition.

Definition 1. Let the home producer's productivity z, the foreign producer's productivity z^*, and wages be given. A Bertrand equilibrium in country j consists of quantities ($q_{hjk}(z, z^*)$,$q_{fjk}(z, z^*)$) and prices $p_{jk}(z, z^*)$ such that

(a) $p_{jk}(z, z^*)$ is equal to the low-cost producer's monopoly price or the high-cost producer's marginal

cost, whichever is lower (see equation 7);

(b) The high-cost producer in a particular location produces a quantity of zero for that location;

(c) The low-cost producer in a particular location produces a quantity for that location given by the

inverse demand function (1).

Three categories of non-traded firms

A firm may fall into one of three categories: *thriving*, *struggling*, or *out*. A firm is an *out* firm if it would not be able to produce profitably even if the manager's efforts to raise productivity were successful. *Out* firms have such low productivities compared to their foreign competitors that they would still be the high-cost producer in each location if their productivities were raised by a factor Δ. *Thriving* firms are those that are the low-cost producers in at least one location even when their managers fail to raise productivity. Thriving firms never shut down.[5] Finally, *struggling* firms are those that are the low-cost producer in at least one location *only* when the manager's efforts to raise productivity are successful. If their managers fail to rise productivity, then they succumb to foreign competition and do not operate.

2.2.5 Contracting problem

Events within a firm unfold as follows. A firm receives its initial productivity draw, which is public knowledge, and is matched with a manager. Managers are risk-neutral. The firm's owners and the manager negotiate a contract specifying payments to each party contingent on profits. Profits can be either high or low depending on whether or not the manager successfully raises productivity. Owners choose the contract that maximizes their own expected payments, subject to a participation constraint for the manager. Because the manager can choose to work as a production laborer, owners must offer the manager a contract that provides expected utility greater than or equal to the utility that can be gained from earning wage w_h. If the owners do not satisfy the manager's participation constraint, the firm shuts down, and the manager returns to the labor force to earn wage w_h. Also, there are limited liability constraints for both parties: neither the owner nor the manager may receive a negative share of profits in any state. Once the owners and the manager agree on a contract, the manager then exerts effort, and productivity gains (if any) are realized. The manager then chooses the quantity of the firm's product to produce for the domestic and foreign markets, and profits are earned and distributed to each party.

The principal-agent friction in the model arises from the fact that effort is not contractible. If firms could make payments contingent on effort, they could always induce managers to exert effort at the first-best level. Instead, payments are contingent on profits, which imperfectly reflect managerial effort because of the probabilistic nature of productivity gains. The moral hazard problem between owners and managers belongs to the class of problems studied in Sappington (1983). With risk-neutral owners and managers, there exists under certain conditions a profit-sharing contract that induces the manager to exert effort at the first-best level, but such a contract would require the manager to receive a negative payment in the low-profit state. As in Sappington,

[5] Unless they cannot satisfy the manager's participation constraint (see below).

the limited liability constraints in my model prevent such contracts from being formed, leading managers to exert effort at an inefficiently low level.[6]

The optimal contracting problem will take a separate form for each category of traded firm - thriving, struggling, and out. For firms that are "out," the problem is trivial: no contract is ever formed because the firm shuts down immediately upon learning its own productivity and the productivity of its competitor. It is as if firms that are out never existed. The contracting problem for a struggling firm will differ from that of a thriving firm in that the managers and owners of a struggling firm take into account the fact that the manager will lose the benefit $\lambda z^{\sigma-1}$ if the manager fails to raise productivity. Note that we can consider producers of non-traded products to be "thriving." Having no foreign competitors, producers of non-traded products are never threatened with the possibility of being forced to shut down if productivity is not raised.

Optimal contract for thriving firms

First consider the optimal contract for thriving firms. Denote by π_H and π_L the values that profits take when the manager raises productivity and when the manager does not raise productivity. We will have $\pi_H > \pi_L$. To keep things simple in this section, the dependence of π_H and π_L on z (and other parameters) is left implicit in the notation.[7] A contract is a pair $s = (s_H, s_L)$, which represents the profit shares to be received by the manager when profits are high and when they are low. Since owners receive the portion of profits left over after the manager is paid, owners receive the payment (1-s_H)π_H when profits are high and (1-s_L)π_L when profits are low.

Managers are risk neutral and, given their contracts, choose their levels of effort to maximize their expected incomes, less the disutility of effort. Equivalently, the manager's problem is to choose the utility-maximizing probability of success ρ^*:

[6] As Harris and Raviv (1979) show, in environments such as the one considered here where owners and managers are both risk-neutral, the moral hazard problem could be eliminated by having the manager become the residual claimant to firm profits after making a lump-sum payment to the owners. However, I assume that due to insitutional constraints or the managers' wealth constraints, such an arrangement is not possible.

[7] Under the optimal contract, the manager will always maximize profits, given the *ex-post* productivity level of the firm, because it is costless for the manager to do so, and because the manager's pay will always be increasing in profits. We therefore can restrict attention to contracts whose payments are specified for only two values of profits, one associated with high productivity Δz, and the other associated with low productivity z.

$$\rho^*(s_H, s_L) = \operatorname*{argmax}_{\rho\in[0,1]} \rho s_H \pi_H + (1-\rho)s_L\pi_L - e(z,\rho) + \lambda z^{\sigma-1}$$

s.t.

$$e(z,\rho) = Az^{\sigma-1}\exp(b\rho)$$

Given the contract (s_H,s_L), the manager's optimal choice of ρ is

$$\rho^*(s_H, s_L) = \Big(\frac{1}{b}\Big)\log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s_H\pi_H - s_L\pi_L)\Big] \tag{2.12}$$

Note that the manager's effort is increasing in the difference between $s_H\pi_H$ and $s_L\pi_L$. As the manager's "performance bonus" for raising productivity increases, effort increases.[8] Given the initial draw of z, the optimal contracting problem faced by the firm owners is as follows:

$$\max_{s_h,s_L} \rho^*(s_H, s_L)[(1-s_H)\pi_H] + (1-\rho^*(s_H, s_L))[(1-s_L)\pi_L]$$

s.t.

$$s_H \in [0,1]$$

$$s_L \in [0,1]$$

$$\rho^*(s_H, s_L)s_H\pi_H + (1-\rho^*(s_H, s_L))s_L\pi_L - e(\rho^*, z) + \lambda z^{\sigma-1} \geq w_h$$

$\rho^*(s_H, s_L)$ is the solution to the manager's problem

Because of the limited-liability constraint on both managers and owners, neither s_H nor s_L can exceed 1 or be less than zero. The third constraint is the manager's participation constraint. Throughout the remainder of the section, I will assume that the participation constraint does not bind. The case in which it does bind is not central to either the paper's theoretical or quantitative results, so exposition of it is relegated to the appendix.

I am now in a position to describe the optimal contract for a thriving firm, which is summarized in the following proposition.

[8] To prevent the exposition from becoming too tedious, I omit discussion of corner solutions and assume that ρ^* lies strictly between 0 and 1.

Proposition 1. *Assume that the manager's participation constraint does not bind. Under the optimal contract for a thriving firm, the following hold:*

(1) $s_L = 0$

(2) $s_H = max\{s^*, 0\}$ *where* s^* *is the solution to*

$$\Big[\frac{\pi_H}{s^*\pi_H - \theta\pi_L}\Big](1 - s^*) = log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s^*\pi_H)\Big] \tag{2.13}$$

(3) The probability of the manager's efforts to raise productivity being successful is

$$\rho^*(s_H, s_L) = \Big(\frac{1}{b}\Big)log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s_H\pi_H)\Big] \tag{2.14}$$

All proofs in this paper are in the appendix. The logic underlying the proposition is fairly straightforward. Owners pay the manager quasi-rents in order to induce the manager to exert effort to raise the firm's productivity so that the owners may enjoy their share of the resulting high profits. Any payment the manager receives in the low-productivity state reduces the manager's incentive to exert effort, so the owners set $s_L = 0$. The owners then choose the manager's payment $s^*\pi_H$ in the high-productivity state so that the expected marginal increase in owners' income from an additional unit of managerial effort is equal to the marginal cost of inducing the manager to exert one more unit of effort.

Under the functional form assumptions I have made for $e(z, \rho)$, it turns out that all producers of non-traded products will have the same probability of raising productivity. That's because $\pi(z) \propto z^{\sigma-1}$ for firms that act as monopolists (see equations 8 and 10), so z drops out of the expression for ρ^*. The model therefore returns a result akin to Gibrat's Law, at least for producers of non-traded products. Because a firm's size is proportional to $z^{\sigma-1}$, all producers of non-traded products will have the same probability of raising their productivity regardless of their initial productivity draws.

Optimal contract for struggling firms

When a firm with initial productivity z is struggling, managers and owners take into account the fact that the manager will lose the benefit $\lambda z^{\sigma-1}$ if the manager cannot raise productivity. With profits in the low-productivity state equal to zero, the manager's problem becomes

$$\rho^*(s_H) = \underset{\rho\in[0,1]}{\text{argmax}}\, \rho s_H \pi_H + \rho\lambda z^{\sigma-1} - e(z,\rho)$$

s.t.

$$e(z,\rho) = Az^{\sigma-1}\exp(b\rho)$$

When profits are zero in the low-productivity state, ρ^* is a function only of s_H. Given s_H, the manager's optimal choice of ρ is

$$\rho^*(s_H) = \Big(\frac{1}{b}\Big)\log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s_H\pi_H + \lambda z^{\sigma-1})\Big] \tag{2.15}$$

As the above expression makes evident, the threat of losing the perks valued at $\lambda z^{\sigma-1}$ serves as an additional incentive for managers to exert a higher level of effort.

Since struggling firms earn no profits when productivity is low, the owners only choose s_H, and a contract s is just $s = s_H$. Their optimal contracting problem becomes

$$\max_{s_H} \rho^*(s_H)[(1-s_H)\pi_H]$$

s.t.

$$s_H \in [0,1]$$

$$\rho^*(s_H)s_H\pi_H - e(\rho^*, z) + \lambda z^{\sigma-1} \geq w_h$$

$\rho^*(s_H)$ is the solution to the manager's problem

Proposition 2 characterizes the optimal contract for a struggling firm.

Proposition 2. *Assume that the manager's participation constraint does not bind. Under the optimal contract for a struggling firm, the following hold:*

(1) $s_H = max\{s^*, 0\}$ *where* s^* *is the solution to*

$$\Big(\frac{\pi_H}{s^*\pi_H + \lambda z^{\sigma-1}}\Big)(1-s^*) = log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s^*\pi_H + \lambda z^{\sigma-1})\Big] \tag{2.16}$$

(2) The probability of the manager's efforts to raise productivity being successful is given by

$$\rho^*(s_H) = \Big(\frac{1}{b}\Big)log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s^*\pi_H + \lambda z^{\sigma-1})\Big] \tag{2.17}$$

By comparing the equations for ρ^* for thriving firms (14) and for struggling firms (17), it would initially seem that the manager of a struggling firm will exert more effort than the manager of a thriving firm, all else equal. That's because the manager of a thriving firm still earns $\lambda z^{\sigma-1}$ even if productivity is low, while the manager of a struggling firm does not. The manager of the struggling firm has a much stronger incentive to work hard to avoid the low-productivity state. However, there is another effect at work, and it is logically possible that the manager of a struggling firm could exert less effort than the manager of a thriving firm, all else equal. The reason is that the owners of a struggling firm recognize that the manager already wants to work hard to avoid losing the benefits valued at $\lambda z^{\sigma-1}$, so they will offer the manager a lower s_H. In extreme cases, the owners may want to offer the manager $s_H = 0$.[9] In the calibrated model, the manager of a firm that goes from being "thriving" to struggling always exerts higher effort, though.

2.3 Equilibrium

Before defining an equilibrium in the model, I first present a few market clearing conditions. From this point forward, I will change the notation slightly and let $\rho^*(z)$ denote the effort exerted by the manager of a non-traded firm with productivity z. Similarly, let $\rho^*(z, z^*)$ denote the value of ρ chosen by the manager of a firm in a traded industry with productivity z that faces a foreign competitor with productivity z^*. The value of ρ^* should be understood to arise from the optimal contract as described in Propositions 1 and 2. Let $\hat{s}(z)$ and $\hat{s}(z, z^*)$ denote the corresponding contracts.

2.3.1 Market clearing

Trade is balanced, so the home country's income equals its expenditure on the final good, $P_h c_h$. Income consists of the sum of production worker wages and aggregate firm profits (before payments to the manager have been subtracted). The market clearing condition for the goods market can be written as

[9] For sufficiently large z, the non-pecuniary perks $\lambda z^{\sigma-1}$ will be high enough so that the participation constraint does not bind even if the manager receives none of the firm's profits.

$$
\begin{aligned}
P_h c_h = w_h(L_h - m_h) &+ M_{hN} \int_0^\infty \Big[(1-\rho^*(z))\pi_{hN}(z) + \rho^*(z)\pi_{hN}(\Delta z)\Big] dG_{hN}(z) \\
&+ M_E \int_0^\infty \int_0^\infty \Big[(1-\rho^*(z,z^*))\pi_{hE}(z,z^*) + \rho^*(z,z^*)\pi_{hE}(\Delta z, z^*)\Big] dG_{fE}(z^*) dG_{hE}(z) \\
&+ M_I \int_0^\infty \int_0^\infty \Big[(1-\rho^*(z,z^*))\pi_{hI}(z,z^*) + \rho^*(z,z^*)\pi_{hI}(\Delta z, z^*)\Big] dG_{fI}(z^*) dG_{hI}(z)
\end{aligned} \tag{2.18}
$$

where m_h denotes the mass of home country firms in operation. Since each firm has one manager, only a mass of workers equal to $L_h - m_H$ earn wages. The integrals represent profits earned in each of the three industries (non-traded = N, export-oriented = E, and import-competing = I), with a firm's profits in the high- and low-productivity states weighted by the manager's probability of raising productivity. We also have the following labor market clearing condition:

$$
\begin{aligned}
L_h - m_h = M_{hN} &\int_0^\infty \Big[(1-\rho^*(z))\ell_{hN}(z) + \rho^*(z)\ell_{hN}(\Delta z)\Big] dG_{hN}(z) \\
&+ M_{hE} \int_0^\infty \int_0^\infty \Big[(1-\rho^*(z,z^*))\ell_{hE}(z,z^*) + \rho^*(z,z^*)\ell_{hE}(\Delta z, z^*)\Big] dG_{fE}(z^*) dG_{hE}(z) \\
&+ M_{hI} \int_0^\infty \int_0^\infty \Big[(1-\rho^*(z,z^*))\ell_{hI}(z,z^*) + \rho^*(z,z^*)\ell_{hI}(\Delta z, z^*)\Big] dG_{fI}(z^*) dG_{hI}(z)
\end{aligned} \tag{2.19}
$$

Because the home country is much smaller than the foreign country, and because $M_{fN} >> M_{fI}$ and $M_{fN} >> M_{fE}$, we can treat the foreign country as a closed economy for the purposes of computing aggregate foreign equilibrium quantities. The foreign goods market clearing condition is

$$
P_f c_f = w_f L_f + M_{fN} \int_0^\infty \pi_{fN}(z^*) dG_{fN}(z^*) \tag{2.20}
$$

The foreign labor market clearing condition is

$$
L_f = M_{fN} \int_0^\infty \ell_{fN}(z^*) dG_{fN}(z^*) \tag{2.21}
$$

2.3.2 Definition of equilibrium

Take foreign production labor to be the numeraire commodity, so that $w_f = 1$. I am now in a position to give a definition of an equilibrium in the model.

Definition 2. An equilibrium consists of (i) prices, quantities, and labor demands for traded products $p_{ijk}(z, z^*)$, $q_{ijk}(z, z^*)$ and $\ell_{ijk}(z, z^*)$; (ii) prices, quantities, and labor demands for non-traded products $p_{iN}(z)$, $q_{iN}(z)$ and $\ell_{iN}(z)$; (iii) a wage w_h; (iv) profit-sharing contracts for home non-traded firms $\hat{s}(z)$; (v) profit-sharing contracts for home traded firms $\hat{s}(z, z^*)$; and (vi) probabilities $\rho^*(z)$ and $\rho^*(z, z^*)$ such that:

(a) Given w_h, $P_h c_h$, $P_f c_f$, and ex-post firm productivity levels, prices and quantities chosen by producers

of traded products constitute a Bertrand equilibrium as defined in Definition 1;

(b) Given w_h, $P_h c_h$, $P_f c_f$, and ex-post firm productivity levels, prices and quantities chosen by producers

of non-traded products solve their profit maximization problem;

(d) Individual firms' labor demands are given by (2) and (3);

(c) Labor market clearing holds in each location (equations 19 and 21);

(d) Product market clearing holds in each location (equations 18 and 20);

(e) Given w_h and $P_h c_h$, $\hat{s}(z)$ is the optimal contract for the producer of a non-traded product with initial

productivity draw z, and $\rho^*(z)$ is the associated probability that the manager's efforts will raise

productivity;

(f) Given w_h, $P_h c_h$, $P_f c_f$, and the foreign competitor's productivity z^*, $\hat{s}(z, z^*)$ is the optimal contract

for the producer of a traded product with initial productivity draw z, and $\rho^*(z, z^*)$ is the associated

probability that the manager's efforts will raise productivity.

2.3.3 Effect of unilateral trade liberalization on firm productivities

A unilateral liberalization can be modeled as a fall in D_{fh}, the iceberg transportation cost faced by foreign producers who sell to the home country. In partial equilibrium, a fall in tariffs will tend to lead to productivity gains. To see the effect that a fall in D_{fh} has on ρ in partial equilibrium, consider a home firm with productivity $z = 1$ that faces a foreign competitor with productivity $z^* = 1.44$. Also let $\Delta = 1.2$. For simplicity, assume $w_h = w_f = 1$ and that D_{hf} is high enough so that we can ignore the possibility that the home firm sells to the foreign market. The home firm will be unable to produce profitably in its home market, even if it raises its productivity, if $\Delta z < z^*/D_{fh}$. As illustrated in Figure 1, this means the home firm will be "out" for $D_{fh} < 1.2$. For values of D_{fh} above 1.44, the home firm will produce in its own market even if it does not raise productivity (since $z > z^*/D_{fh}$). For values of D_{fh} between 1.2 and 1.44, the home firm is struggling.

As the solid line in Figure 1 shows, a fall in D_{fh} leads a firm to raise its productivity substantially if it pushes the firm from the thriving category into the struggling category.

Figure 1: Manager's choice of ρ

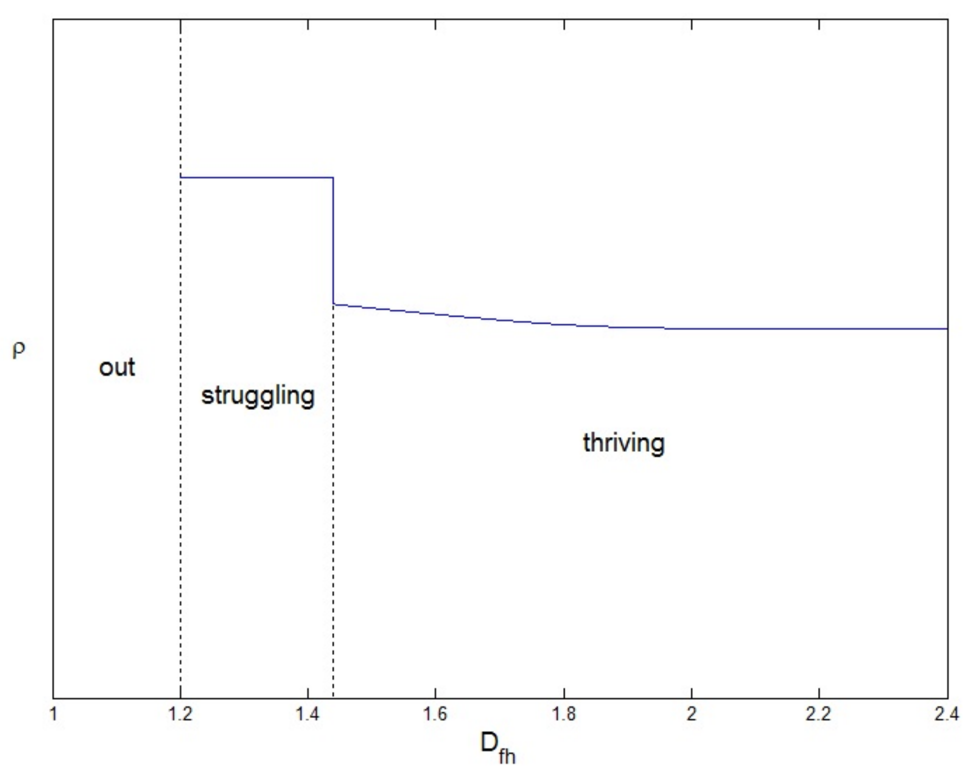

The threat of losing the benefits $\lambda z^{\sigma-1}$ induces the manager to exert more effort. Note also that when a firm is thriving, ρ is initially decreasing for values of D_{fh} close to 1.44 before becoming constant. The reason is that as long as D_{fh} is sufficiently high, the home firm acts as a monopolist (equation 7), and small changes in the foreign competitor's marginal costs do not affect the manager's choice of ρ. For values of D_{fh} closer to but still above 1.44, the home firm acts as a limit-pricer and sets its price equal to the foreign competitor's marginal cost z^*/D_{fh}. As D_{fh} increases, the home firm's price increases, and the quantity it produces falls. When the quantity falls, so does the value of a reduction in unit costs, so owners induce managers to exert less effort to raise productivity.

When general equilibrium effects are incorporated into the analysis, it is possible that a manager will choose a lower value of ρ following a fall in D_{fh}. To see this, consider a fall in D_{fh} from 2.4 to 2.2 so that the home firm acts as a monopolist both before and after the fall in trade costs. Figure 1 suggests that the manager chooses the same value of ρ in each case. However, a fall in trade costs will reduce the residual demand for the firm's product, as some foreign producers of traded products are now able to enter the home market, thereby lowering aggregate prices. When faced with lower demand, the firm produces a lower quantity. With the firm producing a lower quantity, the value of a reduction in unit costs falls. The firm owners then pay the manager less and induce a lower level of effort.

Since the effects of a unilateral fall in trade costs on the productivities of individual firms are ambiguous, the effect on average industry-level productivity will be ambiguous too. Whether or not foreign competitive pressure raises or lowers average firm-level productivity growth within an industry depends critically on the distribution of firm productivities at home and abroad, as well as the magnitude of the parameter λ.

2.3.4 Entrepreneurial firms

I now introduce a second type of firm into the model. Let there be some firms, which I will call *entrepreneurial* firms, that are manager-owned. To this point I have been focusing on firms with separate ownership and management, which I will now call *managerial* firms. In each industry, there are some managerial firms and some entrepreneurial firms. I assume that a firm's ownership structure is given exogenously. Managers of entrepreneurial firms differ only from their counterparts in managerial firms in that they receive all firm profits. Hence their choices of ρ^* are summarized by equations (12) and (15), with $s_H = s_L = 1$.

Qualitatively, the relationship between D_{fh} and ρ^* for an entrepreneurial firm is the same as in Figure 1. However, the productivity "jump" that an entrepreneurial firm takes when it moves from the thriving category to the struggling category is smaller than the jump taken by managerial firms. The reason is that managers of entrepreneurial firms internalize all the fruits of their efforts and therefore are already exerting effort at a high level before a fall in D_{fh} pushes them into the struggling category. Since the disutility is convex in effort (see equation 4), raising ρ incrementally is much more costly for managers of entrepreneurial firms than for managers of managerial firms.

2.4 Evidence from Chile

Between 1974 and 1979, Chile implemented a sweeping program of trade reforms, which included drastically lowering tariffs in the manufacturing sector from an average effective rate of 151 percent to a uniform rate of 10 percent.[10] The result was a large increase in trade volumes in a country that had been fairly closed off to outside goods markets. In 1976, the ratio of trade to output in the manufacturing sector stood at 4.5 percent. By 1986, the figure had risen to 48.5 percent. During the same period, the ratio of imports to output in the manufacturing sector increased from 2.1 percent to 21.7 percent.

In this section, I present two empirical findings about how Chilean plants responded to the increase in foreign competition that resulted from the lower tariffs. Following Liu, Roberts, and Tybout (1996) and Pavcnik (2002), I take 1979-86 to be the period of adjustment to the trade reforms. The year 1978 was the first in which the ratio of trade to output in the manufacturing sector exceeded 10 percent, and the figure remained between 30 and 50 percent in each year during 1978-86. Since plant-level data prior to 1979 are not available, 1979 and not 1978 (or any earlier year) is taken as the initial year of the adjustment period.

Using the same difference in differences regression model as Pavcnik (2002), I first

[10] For more background on Chile's trade reforms, see Edwards and Lederman (1998).

show that continuing plants[11] in import-competing industries saw substantial productivity improvements relative to plants in non-traded industries during the period 1979-86. I then extend the empirical analysis in the simplest possible way to show how productivity growth depends on a firm's ownership type. The plant data allow me to classify businesses based on ownership type. One category, *managerial plants*, consists of plants with separate ownership and management and therefore most afflicted by principal-agent frictions. These are primarily plants in corporations. The second category, *entrepreneurial plants*, consists of those with no separation between ownership and management. The two categories of ownership structure are chosen so that there is a clear mapping between managerial and entrepreneurial firms in the model and in the data. Within the set of managerial plants, I find evidence that plants in import-competing industries have higher average productivity growth than those in non-traded industries. However, within the set of entrepreneurial plants, there is no statistically significant difference in productivity growth among plants in import-competing industries and plants in non-traded industries. In Section 5, I use the estimates of productivity growth differences across plant trade orientations and ownership types to assess how well the calibrated model performs.

2.4.1 Description of the data

Plant-level data come from the census of Chilean manufacturing plants conducted by Chile's National Institute of Statistics. Because of its richness, the census data has been used by numerous researchers (e.g., Liu 1993, Liu, Roberts, and Tybout 1996, and Pavcnik 2002). The census covers all plants with ten or more employees, resulting in a total of 33,394 plant-year observations over the period 1979-86, or an average of 4,770 plant observations per year. In what follows, I restrict attention to the performance of continuing plants that appear in both the first and last years of the sample. Doing so leaves a total of 6,012 plant-year observations in the baseline analysis.

The dataset contains detailed data on firm inputs and outputs, as well as input and output price deflators. Output price deflators are obtained at the 3-digit ISIC level from the Chilean Central Bank's price indices. Firm-specific deflators for intermediate inputs are derived using input-output tables provided by the Chilean Central Bank. (See Liu

[11] defined as plants that appear in both the first and last years of the sample.

1993 for more details.) I construct my key performance measure, labor productivity, for each firm by dividing real output by the average number of workers employed in a given year.[12]

Data on industry-level trade flows come from the World Bank's Trade, Production, and Protection database. The database reports the annual value of industry imports, exports, and output at the 3-digit ISIC level for Chile's manufacturing sector. Using these data, I calculate the ratio of net imports to output, NM/Y, and the ratio of gross trade to output, (X+M)/Y, for each industry, averaged over the period 1979-86. Using a two-stage process, I then classify each industry as belonging to one of three categories: import-competing, export-oriented, or non-traded. First, if the average value of (X+M)/Y is less than 0.15 during 1979-86, I classify the industry as "non-traded." Of the remaining industries, I classify those with $NM/Y > 0.2$ as "import-competing" and those with $NM/Y < 0.2$ as "export-oriented." The cutoff values for NM/Y are chosen so that no industries in the data fall into a potential fourth category in which gross trade volumes are high but net trade is close to zero. All of the traded industries in the data have either $NM/Y < 0.2$ or $NM/Y > 0.2$. There are 11 import-competing industries, 3 export-oriented industries, and 14 non-traded industries. I then assign a trade orientation to each plant according to its 3-digit industry's trade classification. For example, all plants in ISIC industry 381 (fabricated metal products) will be classified as "import-competing plants" because "fabricated metal products" is an import-competing industry. Appendix Table A1 lists the industries by trade orientation.

One issue with the data is the fact that the 3-digit ISIC classification changes for a significant number of plants during the sample period. In some cases (accounting for 3.5 percent of all continuing plants), the plant switches to an industry with a different trade classification. Since the switch to an industry of a different trade orientation may be an endogenous response to the trade reforms, I drop these plants from the set of continuing plants.

I also study differences in plant performance by ownership type. Each plant in the dataset is classified as one of seven business types: sole proprietorship, partnership, cooperative, corporation, collective, public, or other. Among continuing plants, 20 percent are classified as sole proprietorships, 49 percent as partnerships, and 23 percent

[12]Data on employment hours are not available.

Table 1: Classification of business types

Business type	Managerial plant?
Corporation	Yes
Collective	Yes
Cooperative	Yes
Public	Yes
Sole proprietorship	No
Partnership	No

as corporations. The remaining 8 percent are either government-operated businesses, cooperatives, collectives, or "other." The main idea of this paper is that trade competition can lead to firm growth by mitigating principal-agent problems within firms. By exploiting the information on plant ownership types, I am able to provide evidence that the key mechanism in my model is at work during Chile's period of adjustment to trade reforms. Using the information on business types, I create a binary variable called the *Managerial* dummy, which takes on a value of 1 for plants with separate ownership and management. These are the plants with a moral hazard problem and are represented by the managerial plants in the model. Plants that I do not classify as managerial plants are those without separate ownership and management, and they are represented by the entrepreneurial plants in the model. Table 1 shows how I classify plants by ownership structure in the regressions that follow. I omit the category "other" because of its ambiguity.

2.4.2 Productivity growth by trade orientation and ownership type

In this section, I compare labor productivity growth across trade orientations and ownership types during the period 1979-86. To keep terminology consistent with the theoretical section of the paper, industries

are defined by trade orientation, so there are three industries: the import-competing industry, the export-oriented industry, and the non-traded industry. (References to specific ISIC industries, where needed, will be made explicit.) I borrow the following difference in differences framework from Pavcnik (2002) to test for differences in average productivity growth among plants of different trade orientations:

$$\Delta \log(lp_i) = \beta_0 + \beta_T Trade_i + \beta_X X_i + \epsilon_i \quad (2.22)$$

Here, $\Delta\log(lp_i)$ is the log change in labor productivity of plant i from 1979-86. $Trade$ is a vector of indicator variables corresponding to the different trade orientations, and X is a vector of control variables, which in the baseline regression only includes controls for a plant's 2-digit ISIC industry affiliation. (The indicator variable for the non-traded

Table 2: Productivity growth by trade orientation and ownership type

	Dependent Variable: Log change productivity, 1979-86	
	1	2
Import-competing	0.089^b	0.035
	(0.045)	(0.049)
Export-oriented	0.163^a	0.096
	(0.058)	(0.065)
Managerial		0.034
		(0.033)
Import*Managerial		0.131^b
		(0.056)
Export*Managerial		0.244^b
		(0.101)
2-digit ISIC industry dummies	Yes	Yes
Obs	3006	3006

Significance at 1, 5, or 10 % level is denoted by a, b, or c.

industry is the omitted dummy.)

The results are reported in column 1 of Table 2. Import-competing plants have productivity growth 8.9 percentage points higher on average than non-traded plants, and export-oriented plants have productivity growth 16.3 percent higher than their non-traded counterparts. The results for the import-competing industry are roughly in line with Pavcnik (2002), who finds that plants in import-competing industries have average productivity growth 6.9 percentage points higher than plants in non-traded industries. Pavcnik also finds that average productivity growth is 4.9 percentage points higher in the export-oriented industry than the non-traded industry, substantially smaller than the 16.3 percentage point difference I find. The difference in our results can be attributed to differences in our empirical approaches. For example, she uses TFP estimates as her measure of productivity, whereas I use labor productivity.

To ascertain the relationship between ownership type and productivity growth, I modify the specification in (22) by incorporating the *Managerial* dummy and interaction terms between *Managerial* and the trade orientation indicators:

$$\Delta\log(lp_i) = \beta_0 + \beta_1 Managerial_i + \beta_2 Trade + \beta_3 Managerial * Trade + \beta_X X_i + \epsilon_i \quad (2.23)$$

Column 2 of Table 2 presents the results. As it shows, the presence of agency frictions is associated with higher productivity growth, but only for import-competing and export-oriented plants. For non-traded plants, there are no statistically significant differences between plants based on *Managerial* classification. Among plants within the import-competing industry, managerial plants had productivity growth 16.5 percentage points higher than entrepreneurial plants, and among plants in the export-oriented industry, the figure is 27.8.

2.4.3 Robustness

In this section, I address two potential concerns related to the findings in the previous section. First, since I use labor productivity as the key performance measure, it may be that some types of firms experience higher productivity growth than others because they accumulate capital faster during the adjustment period to trade reforms. Using

Table 3: Robustness tests

	Dependent Variable: Log change productivity, 1979-86			
	1	2	3	4
Import-competing	0.093^b	0.035	0.060	0.061
	(0.045)	(0.049)	(0.049)	(0.049)
Export-oriented	0.165^a	0.099	0.088	0.090
	(0.063)	(0.065)	(0.070)	(0.070)
Managerial		0.040	0.039	0.046
		(0.033)	(0.034)	(0.035)
Import*Managerial		0.130^b	0.079	0.077
		(0.058)	(0.060)	(0.060)
Export*Managerial		0.260^b	0.269^b	0.282^b
		(0.102)	(0.108)	(0.109)
materials79		-0.059		-0.067
		(0.038)		(0.046)
k growth	0.057^a		0.057^a	0.057^a
	(0.018)		(0.018)	(0.018)
2-digit ISIC industry dummies	Yes	Yes	Yes	Yes
Obs	2350	3006	2350	2350

Significance at 1, 5, or 10 % level is denoted by a, b, or c.

the subset of firms in my sample for which data on real capital stocks are available, I repeat the regressions from Section 4 with a control for capital growth. I also redo the regression (23) with a control for initial firm size. To the extent that plants classified as managerial plants are larger than those classified as entrepreneurial plants, any productivity growth differences due to size differences may be influencing the results in column 2 of Table 2.

Table 3 presents the results. The variable *k growth* is the log change in a plant's real capital stock between 1979 and 1986. As column 1 shows, differences in capital accumulation are not driving the high labor productivity growth of import-competing and export-oriented plants relative to non-traded plants. While the coefficient on *k growth* is positive (as expected) and significant, the coefficient on the *Import-competing* indicator is actually higher than in column 1 of Table 2 (0.093 versus 0.089). Note though that the sample is different in the two regressions.

Column 2 of Table 3 shows how the relationship between a firm's ownership classification and productivity growth changes when I control for initial plant size. I use the real value of a firm's raw materials inputs in 1979 (*materials79*) as a proxy for size, and as comparison with column 2 of Table 2 shows, including the size control hardly affects the results. For example, without the size control, the coefficient on *Import*Managerial* is 0.131, and with the size control it is 0.130. In each case, the coefficient is significant at the 5% level.[13]

Column 3 shows how controlling for capital growth affects the interaction between a firm's ownership classification and its trade classification. After controlling for capital growth, the coefficient on the interaction term *Import*Managerial* is still positive, but it is smaller than before and no longer significant. This result does not necessarily weaken the central idea of the paper though, which is that foreign competition has a stronger effect on managers' incentives when ownership is separate from management. As mentioned earlier, there is evidence that managers incur private costs associated with making capital investments, and competition could provide them with stronger

[13] In other specifications (not reported), I have used output and employment in 1979 as measures of a plant's initial size, with minimal effect on the results. But since output and employment are by definition correlated with a plant's labor productivity, productivity growth will reflect measurement errors of a plant's initial output and employment. Hence I use raw materials inputs as a proxy for initial size.

incentives to make investments that raise productivity. Column 4 adds the size control to the regression in column 3, and again it does not change the results materially.

2.5 Calibration

I perform a comparative static exercise in which I solve the model with high iceberg transportation costs for foreign producers and compare it to the results when transportation costs are low for foreign producers. I set $D_{fh} = 2.6$ initially to capture the fact that average effective tariffs were 150 percent before Chile's reforms. I also assume an additional 10 percent transportation cost that the government cannot affect, so I have $D_{fh} = 2.6$ instead of $D_{fh} = 2.5$. I then solve the model with $D_{fh} = 1.2$ (representing the uniform 10 percent tariff in manufacturing plus a 10 percent transportation cost). The cost of shipping from the home country to the foreign country is set equal to $D_{hf} = 1.1$ both before and after the unilateral fall in trade costs. The set of continuing firms in the model will be the firms that produce both when D_{fh} is high and when D_{fh} is low. These firms are the focus of my quantitative analysis.

The quantitative strategy, in brief, is to calibrate the model parameters to several targets, including trade flows, and then evaluate how well the model performs in reproducing the two empirical results. When calibrated using data from Chile's trade reform adjustment period 1979-86, the model generates firm productivity growth patterns that are consistent with the two empirical results that were the focus of the previous section. First, import-competing firms in the model have higher growth on average than non-traded firms. Second, in the import-competing industry, managerial firms have higher productivity growth than entrepreneurial firms, but in the non-traded industry the average performance of the two ownership types is roughly the same.

The masses of firms in each industry and of each ownership type are chosen to match their abundance in the data. I assume that each firm in the home country in industry k draws its productivity from an industry-specific lognormal distribution with location parameter μ_{hk} and shape parameter σ^2_{hk}. Foreign firms draw from lognormal distributions governed by parameters μ_{fk} and σ^2_{fk}. I normalize μ_{hN} (corresponding to the non-traded industry) to 1, and the other μ's for the home country are chosen to match the initial mean productivity for each industry in 1979, relative to the non-traded industry. (These and the values of other parameters are presented in Table 4.) I assume values for σ^2_{hN} and each σ^2_{fk}. In particular, I set $\sigma^2_{hN} = 0.1$ and $\sigma^2_{fk} = 0.1$ for each foreign industry k. This parameterization gives productivity dispersions in the home non-traded industry and the foreign industries that are consistent with Syverson (2004).[14] Since I assume σ^2_{fk} is the same for each foreign industry, I denote it henceforth by σ^2_f, dropping the industry subscript.

I take Δ, the fraction by which managers can raise productivity, directly from the data on firm productivity growth. Among the set of continuing firms (in any industry) that had positive productivity growth between 1979 and 1986, the mean rate of productivity growth was 38 percent.[15] I therefore set $\Delta = 1.38$.

There are 9 parameters that remain to be calibrated: σ, A, b, λ, θ, μ_{fN}, μ_{fI}, μ_{fE}, σ^2_{hI} and σ^2_{hE}. Recall that σ is the elasticity of substitution between goods, A is the scaling parameter on the effort function (4), and b is the curvature parameter on the effort function. I arbitrarily set $\mu_{fN} = \mu_{hN}$ ($= 1$), and choosing other reasonable values for μ_{fN} affects the results only negligibly. I will assume $\sigma^2_{hI} = \sigma^2_{hE}$. For the seven remaining parameters, I have seven targets.

Targets for calibration:

1) Ratio of exports to output within the export-oriented industry, 1977

2) Trade/Output ratio within Chile's manufacturing sector in 1977

3) Trade/Output ratio within Chile's manufacturing sector in 1986

4) Ratio of output in import-competing industry to output in non-traded industry, 1977

5) Fraction of import-competing entrepreneurial plants that exit between 1979 and 1986

[14]In principle, one could simply choose the home-country variance parameters to match the productivity dispersions seen in each industry in the Chilean data. However, the Chilean data on plant productivities are noisy and imply large values for σ^2_{hk}. Using these values generates implausible results for a wide range of choices for other parameters: When productivities are highly dispersed, only a small number of highly productive firms in traded industries produce at all.

[15]I drop plants with productivity growth greater than 100 percent to eliminate outlier effects. When all plants with positive productivity growth are included, the mean rate of productivity growth is 53 percent. Measurement error can likely account for the fairly large number of plants that experienced productivity growth above 100 percent.

6) Difference in exit rates between import-competing entrepreneurial plants and import-competing

managerial plants

7) Average markup (computed in the data as net profits divided by gross value of output)

When I calibrate the model, I solve it twice: once with $D_{fh} = 2.6$, and once with $D_{fh} = 1.2$, which correspond to 1979 and 1986, respectively. The year 1979 marks the final year of trade reforms that were put into place beginning in 1974, but it is the first year for which firm productivity data are available. The year 1978 was the first one in which trade/output exceeded 10 percent in the manufacturing sector, so I will assume that firms' productivity responses to the reforms largely took place after 1978. However, precisely because trade volumes jump substantially between 1977 and 1979, it is appropriate to use trade data from 1977 for targets 1-4.

Targets 1-4 help pin down the relative productivities of particular industries in particular locations, and the average markup is most closely related to σ. Target 5 is closely related to the disperson of productivities within the import-competing industry, as captured by the parameter σ^2_{hI}. To see this, suppose the distribution of productivities is degenerate in each industry and that $\mu_{hE} > \mu_{hI}$. Then for certain parameter values, moving from high tariffs to low tariffs could lead to complete specialization as all import-competing firms exit.

The parameter λ is critical to the results, so the strategy to calibrate it also requires some explanation. Target 6 is highly sensitive to the choice of λ and provides the most information about its value of the targets listed above. In the data, 50.1 percent of import-competing entrepreneurial plants exited between 1979 and 1986. Meanwhile, only 41.4 percent of import-competing managerial plants exited, a difference of 8.75 percent between the two categories. As discussed in Section 3, λ has a stronger effect on managerial effort in managerial firms than in entrepreneurial firms. As a result, a higher λ means that fewer import-competing managerial firms exit, relative to import-competing entrepreneurial firms.

Table 4 presents the calibrated parameters. The elasticity of substitution takes on value 3.4, which is somewhat low but is within the range of estimates reported in Broda and Weinstein (2006), who estimate the elasticities of substitution within U.S. industr-

Table 4: Parameter values

Parameter	Description	Value
σ	Elasticity of substitution	3.40
A	Scaling parameter on effort function	0.02
b	Curvature parameter on effort function	3.70
λ	Manager's non-pecuniary benefit	1.50
$(\mu_{hN}, \mu_{hI}, \mu_{hE})$	Domestic prod. mean parameters	(1,0.45,0.85)
$(\mu_{fN}, \mu_{fI}, \mu_{fE})$	Foreign prod. mean parameters	(1,0.81,0.75)
$\sigma^2_{hI} (= \sigma^2_{hE})$	Productivity variance parameter	0.113

ies. To give a sense of scale for λ, the value $\lambda = 1.5$ implies that total non-pecuniary benefits across all managers have value equal to 33 percent of aggregate profits, which at first glance perhaps seems high. However, Dyck and Zingales (2004) show that the private benefits of corporate control can be quite large, especially in developing countries.[16] Moreover, the assumption that managers are risk-neutral inflates the calibrated value of λ: A risk-neutral manager of a struggling firm requires a higher value of λ to exert a given level of effort than does a risk-averse manager of a struggling firm, all else equal.

Table 5 presents the central results of the paper, showing average productivity growth by trade orientation and ownership type, with average productivity growth in the non-traded industry normalized to zero. The data for Table 5 come from columns 1 and 2 in Table 2. In particular, the bold numbers come from column 1, and the differences decomposed by ownership type are derived from column 2. Although not all of the coefficients in Table 2 are statistically significant, they provide a baseline for assessing the model's performance. The model successfully generates the higher productivity growth seen among import-competing firms than among non-traded firms. The average growth among import-competing firms is 6.1 percent higher than among non-traded firms in the model, and 8.9 percent higher in the data. Table 5 also shows differences in behavior between managerial firms and entrepreneurial firms following the

[16]Dyck and Zingales estimate the private benefits associated with having a controlling stake in a corporation in many countries. They find that the private benefits of control are equal in value to 14 percent of a company's value on average, with a maximum value of 65 percent in their data. Estimates of a *manager's* private benefits are harder to come by, but the measured private benefits of corporate control are widely thought to include many of the sorts of perquisites that top executives also enjoy.

tariff reduction. Among import-competing firms in the model, managerial firms have average productivity growth that is 5.6 percent (= 9.6 - 4.0) higher than entrepreneurial firms. This is qualitatively consistent with the data, which show a difference of 16.5 percentage points, but smaller. Among non-traded firms, there is little difference (0.6 percent) in average productivity growth between the two ownership types. In the data, the difference is 3.4 percent, which is both small and statistically insignificant.

Table 5: Average percent productivity growth by industry and business type (continuing plants, relative to non-traded industry)

Industry/Ownership Type	**Model**	**Data**
Import-competing	**6.1**	**8.9**
Managerial	9.6	18.9
Entrepreneurial	4.0	2.4
Non-traded	**0.0**	**0.0**
Managerial	0.4	2.3
Entrepreneurial	-0.2	-1.1

Overall, the model performs well in explaining the high productivity growth of import-competing firms relative to non-traded firms, accounting for 69 percent of the difference in average productivity growth between the two sets of firms. The model also accounts for 34 percent of the difference in average productivity growth between import-competing managerial plants and import-competing entrepreneurial plants. There are many potential reasons why plants of different ownership types might respond differently to increases in foreign competition. For example, corporations may in general have better access to credit or access to better technologies than sole proprietorships. The model shows that the presence of agency frictions is in fact a quantitatively important reason for the differences in productivity growth between the two ownership types.

The model performs less well in matching the data from the export-oriented industry, although the model still replicates many of the qualitative patterns seen in the data for that industry. The export-oriented industry in the model has average productivity growth 4.6 percent higher than the non-traded industry, but in the data the figure is 16.3 percent. Moreover, in the model, the export-oriented industry has lower average productivity growth than the import-competing industry, but the opposite is true in the data. When the productivity gains in the export-oriented industry are decomposed according to the ownership structure of the firms, the story is much the same. While managerial firms in the export-oriented industry have higher average productivity growth than export-oriented entrepreneurial firms (6.6 versus 3.8 percentage points), the difference is not nearly as large as the 27.8 percentage point difference observed in the data. Some possible explanations for the mismatch between the model and the data are given in the next section.

Though the model is successful in matching some key features of relative productivity growth rates across industries and ownership types, it does not match absolute growth rates. In the data, non-traded plants experienced productivity growth of 16.7 percent between 1979 and 1986. In the model, the average non-traded plant's productivity falls by 3.3 percent (though this has been normalized to 0 in Table 4). I do not view this as a significant challenge to the model, though, since the model is not intended to capture trends in productivity growth. Instead, it is designed to capture one-time productivity improvements that take place through the elimination of X-inefficiencies.

2.6 Discussion

The logic that underlies the model results in Table 5 is fairly intuitive. Firms have higher productivity growth when they are struggling, as managers exert more effort to keep the firm profitable and keep their non-pecuniary benefits of working as managers. When tariffs are lowered, foreign competitors have an easier time capturing the domestic market for a product and push firms that were previously "thriving" into the "struggling" category. These firms then experience higher productivity growth. The increased competitive pressure from abroad most affects domestic firms in the import-competing industry because they face highly-productive foreign counterparts who can usurp them when tariffs fall. As a result, the industry experiences substantial productivity growth.[17] Export-oriented firms are subject to the same forces as import-competing

[17] Increased foreign competition will also force some previously "struggling" firms to exit. But these are not continuing firms, which are the only firms considered in Table 5.

firms, but less so. Even with low tariffs, foreign firms have difficulty displacing the relatively produtive export-oriented firms in the domestic market. Compared to the import-competing industry, few firms in the export-oriented industry are forced into the "struggling" category following trade liberalization. So while lower tariffs induce some managers in the export-oriented industry to work harder in order to maintain their non-pecuniary benefits, the effect on average industry productivity growth is smaller than in the import-competing industry.

Non-traded firms, facing no head-to-head foreign competitors, are not directly affected by the trade liberalization. It is only through general equilibrium effects that the liberalization can induce managers of non-traded firms to exert either more or less effort. The net general equilibrium effects turn out to induce slightly negative average productivity growth, as a lower price level P_h (and thus lower demand) more than offsets lower wages: Average productivity among non-traded firms falls by 3.3 percent.

Consider now the fact that in import-competing industries, managerial firms experience higher productivity growth than entrepreneurial firms. In both types of firms, managers will work harder to protect their non-pecuniary benefits when the firm's survival is threatened by a foreign competitor. But prior to the fall in tariffs, the managers of import-competing entrepreneurial firms were *already* working hard to raise productivity. These owner-managers, because they fully internalize the benefits of their own efforts, initially exert substantially more effort than the managers of firms with principal-agent frictions. And because the effort function (4) is convex in ρ, it is harder for managers of entrepreneurial firms to increase ρ further than it is for managers of managerial firms to do so. To put it loosely, managerial firms have more room for improvement. Competition from abroad therefore has a stronger effect on productivity for firms with separate ownership and management.

2.6.1 Real income gains

Real income in the home country of the model is equal to c_h. (Since trade is balanced, real income and real output, c_h, are equal.) In the calibrated model, lowering the marginal trade cost D_{fh} from 2.6 to 1.2 results in real income gains of 1.58 percent. The usual gains arising from comparative advantage contribute to the increase in real income, and further gains come from the endogenous productivity growth of domestic firms. One important source of productivity growth is the reduction of managerial slack in firms that have separate ownership and management. The manager of a "thriving" firm exerts effort below the first-best level if ownership is separate from management, but if foreign competition causes the firm to become a "struggling" firm, then the manager will work harder. The threat of losing the perks $\lambda z^{\sigma-1}$ aligns the incentives of the manager more closely with the interests of the firm owners.

Table 6: Real income growth

Version	% change in real income
Baseline model	1.58
Frictionless model	1.25
Difference	0.33

To quantify the portion of the real income gains attributable to the reduction in principal-agent frictions, I simulate Chile's unilateral liberalization in an alternate, frictionless version of the model. In the frictionless model, all firms are owned by managers. That is, all firms are entrepreneurial firms. The number of firms in each industry is chosen to match the number of all firms in each industry in the data, without regard to ownership type. I use all of the other parameters calibrated in the previous section when solving the frictionless model. As Table 6 shows, when trade costs faced by foreign producers fall from $D_{fh} = 2.6$ to $D_{fh} = 1.2$ in the frictionless model, real income increases by 1.25 percent. Real income gains are 0.33 percentage points lower than in the baseline model, indicating that one cost of trade barriers is to prevent competition from mitigating principal-agent frictions, and that this cost is non-trivial.[18]

2.6.2 Some comments on the export-oriented industry

In the baseline model, productivity growth in the export-oriented industry is not as high as in the data. Also, in the data, export-oriented managerial plants have much higher productivity growth than export-oriented entrepreneurial plants. In the model, the

[18] I am not aware of any other studies that attempt to quantify the losses that result from principal-agent frictions. But as a point of quantitative comparison, Harberger (1954) finds in his classic study that the static deadweight losses from monopoly are equal to around 0.1 percent of output in the U.S.

difference is much smaller. One potential reason for the mismatch between the model and the data is that I have not allowed managers to differ in their skills. A richer model might allow effort requirements to differ across managers. High-skill managers could have a high probability, ρ, of raising productivity while putting in little effort. If the model also included a competitive market for managerial talent, I conjecture that high-skill managers would disproportionately be employed by export-oriented firms, which could potentially amplify the effects of foreign competition on productivity growth in the export-oriented industry.

It is also possible that many of the export-oriented plants in the data are foreign-owned and vertically integrated with their parent companies. Foreign-owned plants may have access to better technology or have stronger incentives to innovate than import-competing firms, which my model does not capture. To the extent that foreign-owned plants are more likely to be classified as corporations, export-oriented managerial plants could be expected to have higher average productivity growth than export-oriented entrepreneurial plants. The data do not contain information on the nationality of a plant's owner, though.

2.7 Conclusion

Firms often raise productivity in response to competitive pressure from abroad following a unilateral fall in trade barriers. This paper focuses on the role of principal-agent frictions between firm owners and managers to explain why a firm might raise productivity in the face of shrinking sales. If competition puts the manager in danger of losing job perks, the manager will work harder to raise productivity. The model considered here can account quantitatively for the relatively high growth seen among plants in import-competing industries following Chile's unilateral tariff reduction in the 1970s, as compared to plants in non-traded goods industries. Moreover, the model accounts for a large portion of the difference in productivity growth between import-competing plants with different ownership structures. The alleviation of principal-agent frictions emerges as a substantial gain from trade that has not previously been quantified.

Chapter 3

Linking Firm Structure and Skill Premium

3.1 Introduction

In this essay we bring new data and new theory to bear on a longstanding question: Why do large firms pay higher wages than smaller firms? While many explanations focus on the fact that large firms employ disproportionately more labor-augmenting technology than smaller firms (e.g. Dunne and Schmitz (1995), Dunne et al. (2002)), we propose a different story: Larger firms choose organizational structures with more layers of management, and the wages paid to managers increase as the number of layers grows. We develop a model in which managers near the top of the management hierarchy in large firms have a larger marginal contribution to output than their counterparts at smaller firms since they oversee a larger quantity of firm activities. As a result, they receive higher wages.

Our model is consistent with empirical evidence that the wage premium offered by large firms is greater among skilled employees than among unskilled employees (e.g. Brown and Medoff (1989), Davis et al. (1991), and Idson and Oi (1999)). Equivalently, the skill premium (the average wage a firm pays to skilled employees divided by the average wage it pays to unskilled employees) is increasing in firm size.

We are able to go a step further, using recent data from Chile's manufacturing sector to show that managers in particular receive a large wage premium at large firms - much more so than other categories of skilled workers. The empirical literature typically uses broad definitions of skilled and unskilled labor to study the skill premium. For example, some studies (e.g. and Pavcnik (2003)) treat blue-collar workers as skilled and white-collar workers as unskilled even though there are substantial skill differences within each group. (For example, both a CEO and a data entry employee would be classified as white-collar.) By using detailed plant-level data on employment and wages from Chile's Manufacturing Industry National Survey (ENIA) during the period 1995-2007, we are able to obtain a more granular picture of the wage premia offered by large firms within different labor categories. We construct and study several alternative measures of the skill premium at the firm level. One measure is the ratio of the average wage a firm's managers receive to the ratio of wages a firm's unskilled production workers receive. Another is the ratio of average wages received by skilled production workers ("technicians") to average wages received by unskilled production workers. The latter ought to most directly capture any skill-bias in the technology hired by large firms.

What we find is striking. The effect of firm size on the manager skill premium is close to 20 times greater than the effect of size on the production worker skill premium. Our results suggest that managerial inputs and firm hierarchies are an important reason why large firms pay a higher skill premium than small firms.

To investigate the relationship between firm size and the skill premium theoretically, we develop a hierarchical model of firms that borrows elements from Caliendo and Rossi-Hansberg (2012) in which large firms optimally choose organizational structures with more layers of management than small firms. A manager receives direct reports from employees on the next-lowest layer (either production workers or lower-ranked managers) and is able to augment their output multiplicatively using her managerial labor. High-level managers at large firms, who have the scope to augment large amounts of output, receive high wages. Small firms, producing relatively little output, have no such high-productivity managers and therefore pay their managers lower wages on average. Additionally, since wages for unskilled workers are constant across firms in our model, the skill premium is increasing in firm size.

An immediate implication of the model is that managers with more direct reports receive higher wages. We find evidence in the Chilean data consistent with this prediction: The skill premium a firm pays to its managers is increasing in the firm's ratio of managers to unskilled workers.

One novel feature of our model is that the number of layers of management a firm chooses is a continuous, rather than discrete variable. Doing so gives the model substantial flexibility and allows us to map the firm's production function into a class of familiar production functions in which workers differ only in the quality units they supply. The advantage of our model, relative to those with just human capital, is that it allows us in addition to evaluate firm hierarchical choices. While the notion of having, say, 2.5 layers of management is not immediately intuitive, we will attempt to provide some clarity later when we present the model.

This paper is connected to an extensive empirical literature on the relationship between firm size and wages. Brown and Medoff (1989) provided an early contribution by showing that large firms tend to pay their workers higher wages than smaller firms. The same phenomenon has been documented by many other researchers. For instance, Davis et al. (1991) study U.S. Census manufacturing data from 1963-86 and find a steep monotonic relationship between a firm's size (in terms of number of employees) and the average wage paid to its workers. Idson and Oi (1999) corroborate their results using more recent Census data.

Most empirical work, including the Davis et al. (1991) study cited above, also shows that the wage gap between large and small firms is the largest among skilled workers, indicating that the within-firm skill premium increases, on average, with firm size. Haskel (1998) provides a counterpoint by showing that, among UK manufacturing firms in the 1980s, small firms had a higher skill premium than larger firms.

Many explanations have been advanced for the wage premium paid by large firms (for a summary, see Katz and Summers (1989)). One explanation for why large firms pay higher wages than small firms that has garnered substantial empirical support is that large firms demand more skilled labor than small firms. For example, Abowd et al. (1999) provide evidence from France, and Zweimuller and Winter-Ebmer (2003) provide evidence from Switzerland. One reason why large firms demand more skilled workers is that they need skilled labor to operate more advanced production technology (See Dunne and Schmitz (1995) or Dunne et al. (2002)). While differences in technology appear to contribute meaningfully to the wage gap between large and small firms, a substantial portion of the variation remains unexplained by technological differences.

The technological explanation for the firm-size wage gap has been adopted by researchers seeking to understand the effects of international trade on labor markets. Bernard and Jensen (1999) document that exporting firms, which tend to be the largest firms, raise the demand for skilled labor relative to unskilled labor, leading to increases in the skill premium. Yeaple (2005) and Bustos (2011) present models explaining exporters' high demand for skilled labor. By serving a large international market, exporters have a stronger incentive to invest in cost-reducing technologies than firms that only serve a small local market. To the extent that the new, labor-saving technologies require skilled labor to operate, large exporting firms will hire more skilled labor than smaller firms. Helpman et al. (2010) further develop this argument theoretically.

Our theoretical approach to studying the firm-size wage gap differs from existing theories in that we focus on the effect of firm size on organizational structure. While our model is consistent with the fact that larger firms demand higher-skilled employees than smaller firms, we develop explicitly the idea that larger firms choose a hierarchical structure with more levels than smaller firms. To do so, we draw on the hierarchical model of Garicano (2000) in which firms optimally choose the number of levels of management by trading off gains from having more managers solve problems against the costs of communication between many levels of management. The basic structure of the model has been adapted to study the effects of international trade on organizational structure (Caliendo and Rossi-Hansberg (2012)) as well as aggregate productivity differences between countries (Grobovsek (2013)).

Relative to their models, our main theoretical contribution is to model the number of management levels that a firm chooses as a continuous, rather than discrete, choice. Doing so demonstrates that our model is essentially a standard, simple human capital model with a restriction on firms which require them to hire specific ratios of workers across levels of education. This gives the ability to make specific predictions about firm organizational structure and skill premia. Instead of being perfectly substitutable in production, imposing a structural hierarchy upon firms turns employees with varying levels of education into complements. Two workers cannot simply replace a manager but must work in conjunction with managers in order to accomplish additional tasks. Larger firms require more levels and a higher average level of human capital, and consequently they pay a higher average wage.

Finally, our paper relates broadly to the extensive literature on the increase in the skill premium over time observed in many countries. Skill-biased technical change is the most prominent (and perhaps the most debated) hypothesis for why the skill premium has risen. While many studies support the skill-biased technical change hypothesis, the evidence is mixed. Card and DiNardo (2002) provide a review of the debate and find the hypothesis lacking. Pavcnik (2003) studies plant-level employment data from Chile - though from an earlier time period than we do - and finds that skill-biased technical change cannot explain the increases in the skill premium observed in her sample. Doms et al. (1997) come to a similar conclusion when studying U.S. manufacturing plant data. In each case, controlling for unobserved plant characteristics eliminates the effect of a plant's technology on the premium it pays to its skilled workers. While we do not focus on the behavior of the skill premium over time, our results resonate with theirs in that technology does not appear to be a key driver of our empirical results concerning the skill premium.

The essay proceeds as follows: In Section 3.2 we present our data findings. In Section 3.3 we introduce a simple, discrete firm structure. Section 3.3 extends the firm structure to allow for a flexible selection of both levels and measure of employees. Section 3.4 presents and characterizes a general equilibrium model with the firm structure introduced in the previous section. Finally, Section 3.5 concludes.

3.2 Data

In this section we use data from the Manufacturing Industry National Survey (ENIA) compiled by the Chilean National Statistics Institute (INE) to analyze the relationship between size and skill premium. This database allows us to differentiate between expenditures on different types of skilled workers. In this way we are able to explain more precisely how size affects skill premium.

ENIA compiles economic and accounting information for registered firms that are established in Chile and have 10 or more employees.[1] We have annual information from 1995 to 2007 for around 4,200 firms per year, although information differentiating types

[1]INE targets firms with 10 or more employees. If a firm has more than one establishment, each establishment is reported separately. In this case, there can be reports by establishments with under 10 employees.

Table 1: Skill Premium by Type of Skilled Worker

Managers	5.56
Technicians	1.70
Owners	2.62
Administrative Personnel	1.41
All skilled workers	1.66
Mean number of workers	48.41
Std. deviation:	75.99

Source: INE. Authors' calculations.

of skilled workers is only available beginning in 2000.

The available information allows us to distinguish between skilled and unskilled workers, as well as how much was spent on employees of each type. Unskilled workers are non specialized workers in charge of executing tasks, mainly manual, that are directly related to the production process ("Blue Collar"). We categorize all other employees as skilled workers ("White Collar").[2]

We now analyze the relation between size and skill premium. We define the skill premium of a firm as the ratio of the average wage per skilled worker to the average wage per unskilled worker. We are able to identify different types of skilled workers in the data. In particular skilled workers can be further divided into owners, managers, technicians and administrative personnel. For each firm we estimate five measures of skill premium: a skill premium for each of the types of skilled workers and a skill premium aggregating across all skilled workers. Table 1 shows the resulting skill premium by type of skilled worker. Note that the average manager gets paid over five and a half times the wage of an unskilled worker. This number is much higher than the skill premium for other types of skilled workers.

Next, we analyze the impact of the size of a firm on the skill premium. To do this we run the following regression:

$$SP_{i,t} = \beta_0 + \beta_1 size_{it} + \alpha_t + \gamma_{ind} + \varepsilon_{it}, \tag{3.1}$$

[2] Only full-time workers are included in the analysis.

where $size_{it}$ is the natural logarithm of the total number of workers at firm i in period t, α_t is a time fixed effect, and γ_{ind} is an industry fixed effect. Table 2 shows that the skill premium is increasing in the size of a firm. This result is robust across all types of skilled workers. Moreover, size seems to affect the skill premium of managers more than it affects the skill premium of other skilled workers. In particular, the effect of size on the skill premium of technicians is very small: Holding all other variables constant, an increase of 1% in the size of a firm raises the skill premium of technicians by 0.0008. That is, even though there is a positive relation, the size of a firm does not seem to be economically significant in determining the skill premium of technicians.

The regression results suggest that bigger firms pay higher wages for skilled workers due to a higher marginal productivity of managers, not of technicians. In particular, the data casts doubt on whether larger firms acquiring skill biased capital can explain why larger firms offer a higher skill premium. [3] Rather, our data analysis supports the importance of considering the organizational structure of a firm when analyzing the relation between size and skill premium.

Table 2: Skill Premium versus Size of Firm

	Managers	Technicians	Owners	Administrative Personnel	All skilled workers
Size	1.767***	0.083***	1.641***	0.130***	0.201***
	(0.031)	(0.006)	(0.033)	(0.005)	(0.004)
Year fixed effects	yes	yes	yes	yes	yes
Industry fixed effects	yes	yes	yes	yes	yes
Number of Observations	13,549	20,381	13,623	22,842	34,291
Adjusted R^2	0.264	0.082	0.204	0.089	0.192

OLS estimation. Standard errors are in parenthesis.

***: Significant at 1%.

Source: INE. Authors' calculations.

We next analyze the relationship between the organizational structure of a firm and the skill premium. For this we extend regression (3.1) to control for the ratio of unskilled workers to managers. We label this ratio as span of control, SC. The extended

[3] We get similar results if we calculate the skill premium using average hourly wage instead of average annual wage.

regression is as follows:

$$SP_{i,t} = \beta_0 + \beta_1 size_{it} + \beta_2 SC_{it} + \alpha_t + \gamma_{ind} + \varepsilon_{it},$$

where α_t and γ_{ind} are defined as in (3.1). Table 3 shows the results. We find that the higher the number of unskilled workers per manager, the larger the difference in wage. This supports the idea that firms where managers have a wider span of control offer a higher compensation to managers.

Table 3: Skill Premium for Managers vs. Composition of Firm

Size	1.619***
	(0.034)
Span of control	0.010***
	(0.001)
Year fixed effects	yes
Industry fixed effects	yes
Number of Observations	13,549
Adjusted R^2	0.267

OLS estimation. Standard errors are in parenthesis.

***: Significant at 1%.

Source: INE. Authors' calculations.

These data findings motivate the structure of firms used in our model. We will focus on the role of managers, abstracting from other types of skilled workers, since it is the wages of managers that drive the relation between size and skill premium. Additionally, the firm structure should explain why managers with more direct reports get paid more.

3.3 Discrete Firm Structure

To develop insight into the relationship between firm size and the skill premium paid to managers, we develop a hierarchical model of firms that draws elements from Garicano (2000) and Caliendo and Rossi-Hansberg (2012). Briefly, a firm consists of several layers of employees, with employees on the bottom layer representing production workers and employees on higher layers representing managers. Production workers complete a number of tasks, but a fraction of their potential output is lost before their tasks are converted into the firm's final good. Managers on the first level above workers receive the uncompleted tasks from the workers who report directly to them and are able to complete a portion of them. Tasks that even a manager on the first level cannot complete are handed to the next layer of management, and so on. A firm chooses the number of layers of management optimally to balance the benefits of being able to solve additional tasks against the costs of compensating the additional managers.

Figure 1(a): Examples of Firm Hierarchies

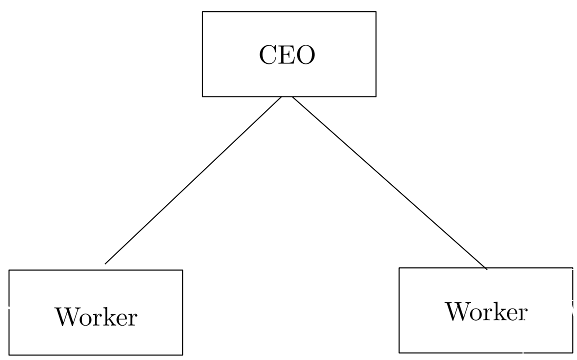

Figure 1(b)

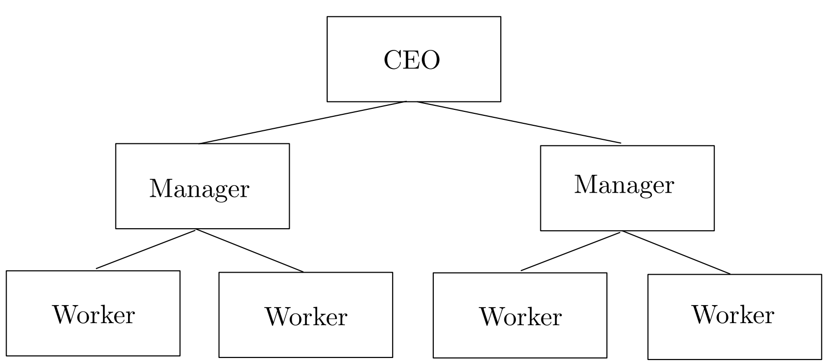

In our full general equilibrium model, the number of management layers chosen by a firm is a continuous variable. We first develop intuition by presenting a simple partial-equilibrium version of the model in which the number of management layers is discrete.

Consider a firm with $X \in \mathbb{N}$ levels. The structure of the firm is rigid in the sense that specific quantities of managers at each level are required. Each worker at level $x = 0$ receives a measure 1 of tasks. Of those tasks, he is able to solve a share equal to $1 - \lambda$, and the uncompleted fraction $\lambda \in (0, 1)$ of the tasks is passed to the level immediately above. No two employees work on the same task. Each manager in level $x > 0$ is assigned R direct reports from the level immediately below and so they receive a measure $R(1 - \lambda)$ of unsolved tasks of which they are able to solve a share $1 - \lambda$. The unsolved fraction λ of the tasks are passed in turn to the next-highest level of management. The quantity R is constant within a firm so that a manager above the first management layer receives direct reports from R managers on the level immediately below. A firm's total output is the number of tasks its workers and managers are collectively able to solve. (In the next section, a firm's output will be a function of the number of tasks completed.)

We refer to R as the span of control of the firm. We denote employees at the lowest level ($x = 0$), who receive no reports, as workers, while employees on higher levels are managers. We will denote managers at the highest level as CEO's and we assume that every firm has one CEO. Figure 1 displays the case where $R = 2$ for firms with $X = 1$ and $X = 2$.

Each firm chooses the optimal number of levels of management X. Hiring additional levels of management allows more tasks to be solved in exchange for a larger wage bill. Specifically, a firm with X levels hires R^{X-x} managers at level $x > 0$ and R^X workers. Furthermore, it is able to solve $R^X(1 - \lambda^{X+1})$ tasks. Lemma 1 states and proves this result.

Lemma 1. *Consider a firm with X levels and span of control R. Each manager in level $x \in \{1, \ldots, X\}$ solves a measure $(R\lambda)^x(1 - \lambda)$ of tasks. The total measure of tasks completed by the firm is $R^X(1 - \lambda^{X+1})$.*

Proof. See B.1. □

Notice that if $R\lambda > 1$ then the contribution of managers is increasing in the level of management x. Assume that wages of all employees - both workers and managers - are proportional to the number of tasks they complete (as they will be in the general equilibrium version of the model). Define the skill premium of a firm, $SP_{X,R}$, as the ratio of the average wage of managers to the average wage of workers. Then $SP_{X,R}$ can be expressed as

$$\begin{aligned} SP_{X,R} &= \frac{\sum_{x=1}^{X} R^{X-x}(R\lambda)^x(1-\lambda)}{\sum_{x=1}^{X} R^{X-x}} \times \frac{1}{1-\lambda} \\ &= \frac{\lambda}{1-\lambda}(R-1)\frac{R^X}{R^X-1}(1-\lambda^X), \end{aligned}$$

where the first term is the average contribution of managers and the second is the inverse of the average contribution of workers. Lemma 2 demonstrates that the skill premium is increasing in both levels of management and span of control.

Lemma 2. *If $R\lambda > 1$ then $SP_{X+1,R} > SP_{X,R}$ and $SP_{X,R+1} > SP_{X,R}$.*

Proof. See B.1. □

The assumption $R\lambda > 1$ implies that each manager completes more tasks than each of her direct reports individually. The lemma establishes two results central to the theme of our paper for the simple, discrete version of the model. The skill premium increases in firm size (proxied by X) because larger firms have top managers who receive many reports (both direct and indirect) and are therefore able to complete many tasks. On average, then, managers at large firms receive higher wages than managers at small firms. The second part of the lemma results from the fact that managers who receive more direct reports have the opportunity to solve more tasks that employees at lower levels were unable to complete.

3.4 Continuous Firm Structure

We now extend the firm structure to allow for a flexible (continuous) selection of both levels and measure of employees. Consider a firm with X levels and span of control R.

To maintain consistency between the discrete version of the model and the continuous version, we assume a width one of CEOs. The span of control R determines the width of managers at level x to be R^{X-x}. Let manager $n \in [0, R^{X-x}-1]$ in level x be a manager who occupies a space of width one with lower edge equal to n. We define the "area of command" for manager n located on level $x > 0$ to be the set of workers and managers who either report directly to manager n or report to a manager within the area of command for manager n. Manager n is in charge of dealing with the unsolved tasks of everyone within her area of command. Formally, the area of command of manager n on level $x > 0$, $AC(x, n)$, is defined as

$$AC(x,n) \equiv \left\{ (x', n') : x' \in [0, x), nR^{x-x'} < n' \leq (n+1)R^{x-x'} \right\}.$$

To map the concept of area of command to the discrete case, consider a firm with 2 levels and span of control $R = 3$. The area of command of manager $n = 1$ on level $x = 1$ is all workers at the bottom level located strictly above position 3 ($= nR^x$) and below position 6 ($= (n+1)R^x$). Figure 2 illustrates an area of command. Lemma 3 characterizes the measure of employees in $AC(x, n)$.

Lemma 3. *There is a continuum of measure* $\frac{1}{\ln R}(R^x - 1)$ *of employees in area of command* $AC(x, n)$*.*

Proof. The measure of managers in $AC(x, n)$ is given by

$$\int_{(x',n') \in AC(x,n)} dx' dn' = \frac{1}{\ln R}(R^x - 1)$$

□

Corollary 1. *Assume that there is a measure 1 of CEOs. Then there is a measure* $\frac{R^X - 1}{\ln R}$ *of managers in a firm with* X *levels and span of control* R*.*

Notice that the measure of employees in an area of command doesn't depend on the location of manager n on level x. Therefore, without loss of generality, we'll denote areas of command and the tasks solved within them by the level of the manager. That is, the area of command of a manager in floor x will be denoted by $AC(x)$.

Similar to before, every employee is able to solve a measure $1 - \lambda$ of the tasks she is assigned. The remaining fraction is sent to a manager located in a higher level. As before, the measure of tasks each employee is able to solve is independent of the tasks any other employee can solve. Now, the number of reports that a manager in level x gets from employees in level $x - \Delta$ is equal to R^{Δ}. Intuitively, as we transition from the discrete to continuous case, levels come closer together and each manager gets less reports from the level directly below them, but there are more levels below them. Additionally, they are able to solve only a fraction $(1 - \lambda)\Delta$ of the tasks they receive. Proposition 3 characterizes the number of tasks that are solved in the area of command $AC(x)$.

Figure 2: Area of Command

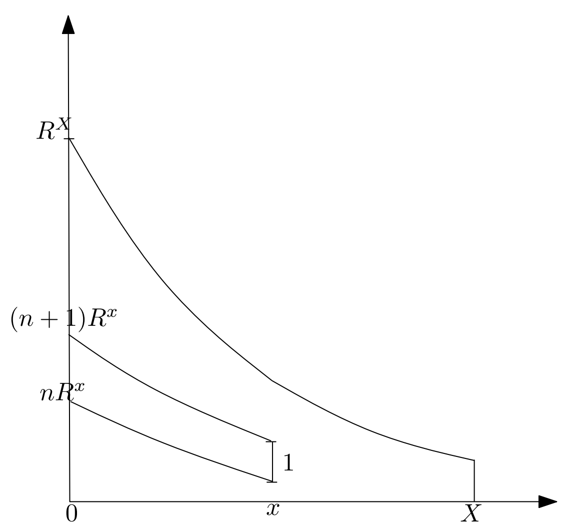

Proposition 3. *Suppose each worker receives a measure one of tasks. Let $\mathcal{T}(x)$ be the number of tasks that are solved in the area of command $AC(x)$. Then*

$$\mathcal{T}(x) = R^x \left(1 - e^{-x(1-\lambda)}\right).$$

Proof. See B.1. □

Corollary 2. *Assume that there is a measure 1 of CEOs. Then the total number of tasks solved by a firm with X levels and R reports is $R^X(1 - e^{-X(1-\lambda)})$.*

Now, the area of command of a manager in level x, $AC(x)$, receives a measure R^x of tasks to be solved. From Proposition 3, the total measure of tasks solved within this area of command is $R^x\left(1-e^{-x(1-\lambda)}\right)$, so there is a measure of $R^x - R^x\left(1-e^{-x(1-\lambda)}\right) = R^x e^{-x(1-\lambda)}$ of unsolved tasks that a manager at floor x faces. Of this measure she is able to solve a fraction $1-\lambda$. Therefore the marginal contribution of a manager at level x, MC_x is given by

$$MC_x \equiv R^x e^{-x(1-\lambda)}(1-\lambda). \tag{3.2}$$

3.5 General Equilibrium Model

We now embed the structure introduced in Section 3.4 in a general equilibrium framework with consumers, firms and an education sector. Consumers supply labor inelastically and choose a quantity of education. The education a consumer receives determines the level of a firm the consumer is qualified to work at and the wage he receives. Consumers use their labor income to pay for education and the consumption good. Only by receiving education can an individual work as a manager. Firms operate a technology that converts completed tasks into the final good with heterogeneous productivities, and they choose an optimal number of levels of management, which in our model is equivalent to choosing the size of the firm. The education sector employs educators and is competitive.

3.5.1 Consumers

There is a measure L of consumers. All consumers are identical at the beginning of each period and own an equal share of all firms. Consumers choose whether to work in the education sector or the private sector by comparing the utility received from working in the education sector, V_E, with the utility received from working in the private sector, V_P.

In order to obtain competence in educating others, educators require $h_E < 1$ units of instruction from other educators. Educators supply instruction in the education sector inelastically and receive a wage of w_E. The budget constraint for an educator is

$$c_E + w_E h_E = w_E + \frac{\Pi}{L}.$$

This implies that the utility an educator receives is

$$V_E \equiv u\left(w_E(1-h_E)+\frac{\Pi}{L}\right),$$

where Π denotes the total profits across all firms and $u(\cdot)$ is the utility function.

Consumers in the private sector choose a level of education x. The quantity x also signifies the level at which the consumer is able to manage at a firm. For example, an individual with a level of education $x = 1$ can work at the first (lowest) level of management. The amount of time required for an educator to train a consumer to work at level x is denoted by $h(x)$. Formally, the problem of a consumer in the private sector is:

$$\begin{aligned} V_P &\equiv \max_x u(c) \\ \text{s.t. } & c + w_E h(x) = w(x) + \frac{\Pi}{L} \end{aligned} \tag{3.3}$$

Since all consumers are ex-ante identical, it must be the case that consumers are indifferent between all choices in their optimal set. If we assume that there exists an x_E such that the level of education required to work at level x_E is equivalent to the level of education required to become an educator (i.e. $h(x_E) = h_E$), then the indifference requirement immediately implies that $w_E = w(x_E)$. This allows us to simplify the consumer's problem to be (3.3) since V_E is given by the utility obtained from choosing x_E.

Notice that in order to educate a manager to level x we require $h(x)$ units of educator time. Further, each of those educators requires h_E units of educator time, and those educators also need h_E units of educators time, and so on. Therefore the total time required in order to train a manager to level x is

$$h(x) + h(x)h_E + (h(x)h_E)h_E + \cdots = h(x)\sum_{n=0}^{\infty} h_E^n = \frac{h(x)}{1-h_E}.$$

3.5.2 Firms

The homogeneous final good is produced by a continuum of firms. A firm whose employees complete $\mathcal{T}$ total tasks produces the final good using technology $Y(\mathcal{T}, z)$. We make basic regularity assumptions on Y; namely, $Y_{\mathcal{T}} > 0$, $Y_{\mathcal{T}\mathcal{T}} < 0$, $Y_z > 0$, and $Y(\mathcal{T}, z)$ satisfies Inada conditions on $\mathcal{T}$. Firms have idiosyncratic technology parameters z distributed according to the density function $f(\cdot)$. Given z, the firm chooses a number of employee levels X in order to maximize profits. The structure of the firm is rigid in the sense that a firm with X levels must employ R^{X-x} employees at each level $x \in [0, X]$ at the firm. In this sense, the firm faces a modified Leontief production function for which the proportions of employees hired at each level are fixed. Workers and managers are able to complete tasks as outlined in Section 3.4. Let $\mathcal{T}(X)$ denote the number of tasks completed by a firm with X levels.[4] Profits for firms are given by

$$\begin{aligned} \pi(z) = &\max_X Y(\mathcal{T}(X), z) - \int_x w(x) l(x, X) dx \\ &\text{s.t. } l(x, X) = R^{X-x}. \end{aligned} \tag{3.4}$$

Total profits across all firms are then $\Pi = \int_z \pi(z) f(z) dz$.

3.5.3 Relation to the Standard Model with Human Capital

At its core, our model is a model with human capital or education with the added restriction that firms hire specific ratios of workers across levels of education. Instead of being perfectly substitutable in production, imposing a structural hierarchy upon firms turns employees with varying levels of education into complements. Two workers cannot simply replace a manager but must work in conjunction with managers in order to accomplish additional tasks. Larger firms require more layers and ultimately a higher average level of human capital. Consequently, they pay a higher average wage.

3.5.4 Equilibrium

An equilibrium is a function which maps education levels to wages, $w : \mathcal{X} \to \mathbb{R}_+$, and a function which maps firm technology endowments to an optimal number of levels for

[4] Recall that $\mathcal{T}(X) = R^X \left(1 - e^{-X(1-\lambda)}\right)$; see the Corollary to Proposition 3.

the firm, $X^* : \mathcal{Z} \to \mathbb{R}_+$, such that:

1. X^* solves the firm's problem (3.4);

2. consumers are indifferent between any education level which is employed with positive measure;

3. the total labor market clears

$$\int_z \int_x l(x, X^*(z)) \left(1 + \frac{h(x)}{1 - h_E}\right) f(z) dx dz = L;$$

4. the goods market clears

$$\int_z Y(\mathcal{T}(X^*(z)), z) f(z) dz = \int_{x \neq x_E} c(x) g(x) dx + M_E c(x_E),$$

where

$$M_E \equiv \int_z \int_x l(x, X^*(z)) \frac{h(x)}{1 - h_E} f(z) dx dz$$
$$g(x) \equiv \frac{1}{L} \int_z l(x, X^*(z)) f(z) dz.$$

3.5.5 Solving the Equilibrium

We make assumptions about the functional form of $h(x)$ and the firm structure in order to obtain an analytical solution to the model.

Assumption 1.

$$h(x) = (R^x e^{-x(1-\lambda)} - 1)(1 - h_E)$$

Assumption 2. *The span of control is sufficiently wide, specifically:*

$$\ln R > 1 - \lambda$$

We denote by w_0 the wage rate for an unskilled worker on the bottom level of a firm, who has received no education. By equilibrium condition 2 above we know that consumers must be indifferent between obtaining any education level which is employed with positive measure. This implies that $w(x) = w_E h(x) + w_0$. Assumption 1, the marginal productivity of a worker at level x from equation (3.2) and the fact that $w_E = \frac{w_0}{1-h_E}$ imply that the marginal wage bill for completing an additional task is constant and equal to $\frac{w_0}{1-\lambda}$. Therefore w_0 is sufficient for characterizing the entire wage schedule.

Now, due to the rigid nature of the firm's structure and Assumption 2 on the span of control R of the firm, the mapping from X to $\mathcal{T}$ is monotonic and we can simplify the firm's problem to choosing the optimal number of tasks to be solved. Proposition 4 formally states and proves this.

Proposition 4. *Under Assumptions 1 and 2 the firm's problem becomes:*

$$\max_{\mathcal{T}} Y(\mathcal{T}, z) - w_0 \mathcal{T}. \tag{3.5}$$

Proof. See Appendix. □

This immediately implies that a firm with a higher technology parameter z will solve more tasks. Since tasks are monotonically increasing in the number of levels, it is also straightforward that the number of levels in a firm is increasing in z. The continuous case offers an analytical advantage over the discrete case since solutions do not depend on cutoff values for z. The continuity in selection of tasks along z allows us to demonstrate the existence of the equilibrium which we formalize with Proposition 5.

Proposition 5. *Suppose Assumptions 1 and 2 hold. Then an equilibrium exists.*

Proof. See B.1. □

Corollary 3. *If $Y(\mathcal{T}, z) = z\mathcal{T}^{\alpha}$, then*

$$\mathcal{T}^*(z) = \left(\frac{\alpha z}{w_0}\right)^{\frac{1}{1-\alpha}}$$

for each z and

$$w_0 = \alpha \left(\frac{E\left[z^{\frac{1}{1-\alpha}}\right]}{L(1-\lambda)}\right)^{1-\alpha}$$

clears the labor market.

3.5.6 Skill Premium

The model can help us understand the effects of firm size and structure on skill premium. We first develop the notion of a skill premium within the model. Similar to section 3.3, we define the skill premium for a firm with X levels and span of control R to be the ratio of the average wage for managers to the wage for workers; that is,

$$SP(X,R) \equiv \frac{\int_{x>0} w(x)l(x,X)dx / \int_{x>0} l(x,X)dx}{w_0}.$$

This aligns well with the definition we use in the data where skill premium is defined as the ratio of average wages for managers to average wages of blue collar workers. Using this definition of skill premium, we are able to show that skill premium is increasing in the size of the firm.

Lemma 4. *Under Assumptions 1 and 2,* $\frac{\partial SP}{\partial X} > 0$.

Proof. See B.1. □

Next, we show that the skill premium is increasing in the span of control of the firm.

Lemma 5. *Under Assumption 1,* $\frac{\partial SP}{\partial R} > 0$.

Proof. See B.1. □

In our model, there is a monotonic relationship between the size of firms and the number of levels the firm employs. Therefore, Lemma 4 aligns with our first empirical observation, namely, that skill premium is increasing with size. This is a direct result of the fact that managers at higher levels have a higher marginal productivity than managers at lower levels and workers. As more levels are added, the average productivity of managers increases and since the productivity per worker is constant in size, the result is obtained.

The intuition for our second result is similar. The number of tasks passing through a manager in a given level is increasing in the number of his direct reports. As the manager's span of control increases so does his marginal contribution to the firm. In our model this directly correlates to the manager's compensation due to our educational structure, thus our second result 5. This result can be generalized by relaxing our assumption on education and adding consumer heterogeneity. In any model where marginal contribution and wages are positively related a similar result will hold. This is also consistent with the second fact we observe in the data, namely, that skill premium is increasing in the ratio of workers per manager.

3.6 Conclusion

In this essay we use recent data from the Chilean Manufacturing Survey to document that, consistent with previous findings, skill premium is positively related with the size of firms. We exploit the rich information we have on different types of skilled workers to estimate the relationship between skill premia by worker type and size. We find that the effect of size on skill premium is much greater for managers than for technicians. From this we conclude that the organizational structure of the firm is important in explaining the positive relationship between size and skill premium. Further findings on the positive correlation between the ratio of workers to managers and skill premium suggest that span of control plays a key role in a firm's organizational structure.

We build on the discrete structural models in the literature à la Garicano (2000) and Caliendo and Rossi-Hansberg (2012) by developing a simple continuous version of the model which maps higher levels of management to higher levels of human capital. The model allows for a closed form analytical solution and provides comparative statics which make specific predictions about the affects of organizational structure on its wage structure; namely, the two facts in the data outlined above.

Chapter 4

Market Structure, Innovation, and Allocative Efficiency

4.1 Introduction

In this essay, I develop a model to study how firm technological upgrading decisions affect allocative efficiency in a multi-industry environment where market concentration differs across industries. I find that for a range of empirically relevant numerical simulations, technological upgrading increases the allocative inefficiency that results from differences in market power across industries. Firms in industries with a low number of producers and high market concentration do not upgrade their productive technology, while firms in more competitive industries do. As a result, resources shift away from industries that are already producing too few goods relative to the socially optimal level. Allocative inefficiency increases. The main result of the essay, then, is once that firm upgrading decisions are taken into account, the efficiency losses from differences in market concentration become even larger.

Empirical work suggests that there is an "inverse-U" relationship between market concentration and firm innovations (Aghion et al. 2009). Competition promotes innovative efforts, but only up to a point: Firms in highly competitive industries have little to gain from productivity improvements since there are few rents to be had. To my knowledge, the model developed here is the first to capture the inverse-U relationship in a general equilibrium setting. The advantage of developing a general equilibrium model is that it allows me to perform welfare analysis and study how firm innovation decisions affect allocative efficiency. Reksulak et al (2008) come to the conclusion that innovations made by monopolists reduce welfare, arguing that when a monopolist raises productivity, the difference between a monopoly's actual output and the socially optimal level increases. My model implies the opposite, which is that welfare ought to increase when a monopoly raises its productivity. The difference between my paper and theirs largely arises from the fact that their analysis is partial equilibrium, and the fact that I find a result that is the opposite of theirs underscores the importance of studying welfare in a general equilibrium setting. The paper most similar to mine is Aghion et al (2009). The authors develop an inverse-U model of competition and innovation in a partial equilibrium environment, but they do not seek to study welfare.

The model I study captures, broadly, the same economic forces that govern innovation decisions in the Aghion et al. model. In my model, there are many industries, and the exogenously-given number of producers varies across industries. Each industry has a leading firm that is able to upgrade its technology and produce at a lower unit cost than its competitors. The inverse-U relationship arises from two countervailing forces. First, as the number of producers in an industry increases, rents for each producer fall as each producers market share shrinks and the price is driven down closer to firms marginal costs. When rents are low, the benefit of raising productivity falls, all else equal, leading to less technological upgrading. Competition also has a positive effect on upgrading, though. In highly competitive markets where firms each have small market shares, the potential gains from having a productivity advantage over rival firms is potentially large. By raising productivity, one firm can improve its market share substantially, which represents a strong incentive to upgrade. This is similar to the "escape from competition" motive present in Aghion et al. Monopolists, by contrast, do not face the same incentive because they capture the entire market whether or not they upgrade. When the two forces combine, there is an inverse-U relationship between market concentration and innovation.

The inverse-U relationship represents a middle ground between the Schumpeterian and Arrow views of competition and innovation. The Schumpeterian view (Schumpeter 1942) emphasizes the need for a firm to be able to earn rents in order to have an incentive to invest in technological improvements. Firms with substantial market power should be driving innovation, then. The Arrow view (Arrow 1962), which has comparatively more evidence in its favor[1], is that product market competition fosters innovation. Under certain parameterizations, a firms gains from upgrading are strictly increasing in the competitiveness of an industry (measured by the number of producers in the industry). While I do not focus on this case, the effects of upgrading on allocative efficiency are somewhat stronger than in the baseline inverse-U case.

The question of how firms' innovations affect allocative inefficiency is not just of theoretical interest. According to Sidak and Teece (2009), "[F]ederal courts have...caused antitrust case law to ossify around a decidedly static view of antitrust." Many, including the authors, have argued that antitrust law should take a more dynamic view of the competitive environment and consider how breaking up a monopoly would affect innovation in a particular industry. The discussion is typically framed as a question of weighing static allocative efficiency against dynamic technical efficiency. That is, monopoly may result in allocative inefficiency today, but investments in innovation made by monopolists lead to productivity that is closer to the socially optimal level. (These studies - mostly but not always - seem to lean toward a Schumpeterian view of innovation in which firms with extensive market power are making the major innovations.) The message of my essay is that today's market structure also affects allocative efficiency in the future through its effect on firm innovations. A better understanding of the linkages between competition, innovation, and efficiency can inform discussions about public interest in the context of antitrust law.

The remainder of the essay proceeds as follows. Section 4.2 presents the equilibrium model. Section 4.3 presents the corresponding social planner problem. In Section 4.4, I conduct welfare analysis and compare the welfare losses in my baseline model to a version of the model with no upgrading. I show that the welfare losses stemming from allocative inefficiency are larger in the baseline model for a range of parameterizations. Section 4.5 concludes.

[1]See Holmes and Schmitz (2010) for a review of the evidence.

4.2 Model

I develop a multi-industry general equilibrium model in which industries are distinguished by their competitiveness (in terms of the number of producers). In each industry there is a leader who can push the industry's technological frontier outward by paying a fixed cost, along with potentially many other "follower" producers. Whether or not the leader upgrades depends on the number of other producers in the industry. The pattern of upgrading exhibits an inverse-U pattern in the sense that leaders in industries with intermediate levels of competition have the strongest incentives to upgrade.

4.2.1 Preferences

There is a representative agent of unit mass who has linear preferences over a final good c. The final good is a CES aggregate of differentiated products indexed by j:

$$c = \Big(\int_0^1 q(j)^{\frac{\sigma-1}{\sigma}} dj \Big)^{\frac{\sigma}{\sigma-1}}$$

where σ is the elasticity of substitution between products. By standard arguments, the inverse demand for a particular product is given by:

$$p(j) = P\Big(\frac{c}{q(j)} \Big)^{\frac{1}{\sigma}} \tag{4.1}$$

where $p(j)$ is the price of product j, $q(j)$ is the quantity of the product, and P is the ideal price index defined as:

$$P \equiv \Big(\int_0^1 p(j)^{1-\sigma} dj \Big)^{\frac{1}{1-\sigma}}$$

4.2.2 Production

In each industry j there are $n(j)$ firms, $n(j) \in [1, \infty)$. Since industries are only distinguished by the number of competitors in them, I will from now on index an industry by n instead of j. Preferences can therefore be rewritten as:

$$c = \Big(\int_1^\infty q(n)^{\frac{\sigma-1}{\sigma}} df(n) \Big)^{\frac{\sigma}{\sigma-1}}$$

where $f(n)$ is the desnity of industries with n producers.

Each firm produces an identical version of product n, and firms are Cournot competitors. The number of producers in an industry is given exogenously, and making n endogenous (for example, by having firms pay a fixed entry cost) would not substantially change the results. Firms use a linear labor technology with idiosyncratic productivity $z_i(n)$, where i indexes an individual firm in industry n. The representative agent supplies one unit of aggregate labor inelastically. I designate labor as the numeraire commdity, so the wage is equal to 1 throughout. In each industry, there is a single "leader" that produces with productivity γz and $n-1$ "followers" that produce with productivity z, where $\gamma > 1$. I assume z is constant across industries. I will generally - though not always - carry the term z throughout the exposition, but I set $z = 1$ without loss of generality.

Before production takes place, followers can attempt (costlessly) to imitate the leader's technology. With probablility p, they succeed. For now I focus on the case $p = 1$, so that all followers can produce with productivity γz. What distinguishes the leader from the followers is that the leader has the opportunity to upgrade its productivity to $\gamma^2 z$ by paying a fixed cost f, denominated in units of the final good. Followers cannot imitate the new, upgraded technology. The leader will upgrade if

$$\pi^I(n) - f \geq \pi^N(n)$$

where $\pi^I(n)$ are the leader's variable profits if it upgrades and $\pi^N(n)$ are the leader's variable profits if it does not.

4.2.3 Profit maximization problems

If the leader does not upgrade, all n firms in the industry produce with productivity γ. Each producer's profit maximization problem is

$$\begin{gathered}
\text{Max } p(n)q_i(n) - \frac{q_i(n)}{\gamma z} \\
\text{s.t.} \\
p(n) = Pc^{\frac{1}{\sigma}}\Big(\sum_{i=1}^{n} q_i(n)\Big)^{-\frac{1}{\sigma}} \\
q_i(n) \geq 0 \\
q_{-i}(n) \text{ given}
\end{gathered}$$

Throughout, I will focus on symmetric Cournot equilibria when all producers in an industry have the same productivity. Solving for the equilibrium is routine, so I omit the details. The price in industry n when the leader does not upgrade is given by:

$$p(n) = \Big(\frac{\sigma n}{\sigma n - 1}\Big)\Big(\frac{1}{\gamma z}\Big) \tag{4.2}$$

When it does not upgrade, the leader's profits, $\pi^N(n)$, are the same as profits earned by all other firms in the industry. They are given by:

$$\pi^N(n) = P^\sigma c\Big(\frac{1}{n}\Big)\Big(\frac{1}{\sigma n - 1}\Big)(\gamma z)^{\sigma-1} \tag{4.3}$$

Note that when $n = 1$, the price is the familiar monopoly price, which is a constant markup $\sigma/(\sigma - 1)$ times marginal cost. Profits are also equal to the corresponding monopoly profits. Meanwhile, as $n \to \infty$ (which corresponds to the perfectly competitive case), the price approaches firms' marginal cost and profits approach zero.

When the leader upgrades, the problem is similar, but now the leader produces with productivity γ^2 while all firms have productivity γ. A follower's profit-maximization problem is the same as the one above, and the leader's profits maximization problem is:

$$\text{Max } p(n)q_i(n) - \frac{q_i(n)}{\gamma^2 z}$$
$$\text{s.t.}$$
$$p(n) = Pc^{\frac{1}{\sigma}}\Big(\sum_{i=1}^{n} q_i(n)\Big)^{-\frac{1}{\sigma}}$$

Since firms are no longer symmetric, the equilibrium will not be symmetric. Instead, I will consider equilibria in which all followers produce the same quantity. The leader will produce a larger quantity. There are two cases to consider. First, when $\gamma \geq \sigma/(\sigma - 1)$, the leader captures the entire market and the followers produce nothing. What this condition says is that the leader's productivity advantage over its competitors (the left side of the inequality) is greater than the monopoly markup (the right side of the inequality). When the inequality holds, the leader is able to set a price equal to the monopoly markup over its own marginal cost, and that price will still be below the followers' marginal costs. Hence the followers produce zero. When $\gamma < \sigma/(\sigma - 1)$, all producers in the industry produce a positive quantity. The expressions below summarize the price charged in the industry.

$$p(n) = \begin{cases} \left(\frac{\sigma[\gamma(n-1)+1]}{\sigma n - 1}\right)\left(\frac{1}{\gamma^2 z}\right), & \text{if } \gamma < \sigma/(\sigma-1). \\ \left(\frac{\sigma}{\sigma-1}\right)\left(\frac{1}{\gamma^2 z}\right), & \text{otherwise.} \end{cases} \tag{4.4}$$

For the case $n = 1$, note that the leader charges the monopoly price in either case. For the perfectly competitive case $n \to \infty$ in which followers produce non-zero quantities, the price converges to the followers' marginal cost:

$$\lim_{n\to\infty} p(n) = \frac{1}{\gamma z}$$

In this case, the difference between the followers' marginal cost and the leader's marginal cost represents the leader's per-unit profit.

4.2.4 Upgrading decisions

The difference between $\pi^I(n)$ and $\pi^N(n)$ measures how strong the leader's incentive to upgrade is. In Figure 1, I plot $\pi^I(n) - \pi^N(n)$ as a function of n. The parameters used to create the graph (see Table 1) were not chosen in any systematic way, but many reasonable parameterizations produce a figure that looks qualitatively similar.

As the figure shows, there is an inverse-U relationship between market concentration and a leader's incentive to upgrade. Depending on the value of f, there are three outcomes. For f sufficiently low, every leader upgrades. For f sufficiently high, no one upgrades. For intermediate f, we see more economically interesting behavior. Figure 1 illustrates innovative behavior for two different fixed costs, f_1 and f_2. For the higher of the two costs, f_1, only leaders facing an intermediate amount of competition choose to upgrade. Neither the monopolists nor the leaders in highly competitive industries have a strong enough incentive to pay the fixed cost.

Table 1: Parameterization for Figure 1

Parameter	Value
σ	3.5
γ	1.2
z	1

There is another economically interesting case when $f = f_2$. In it, f_2 is such that leaders in all but the least competitive industries upgrade. One can show for this case that as n approaches infinity, the difference $\pi^I(n) - \pi^N(n)$ decreases asymptotically toward a value greater than $\pi^I(1) - \pi^N(1)$. The level of f_2 depicted in Figure 1 could be thought of as capturing the "Arrow view" of competition and innovation, in which the industries with the most innovation are those that are the most competitive. Note,

Figure 1: Gains from upgrading

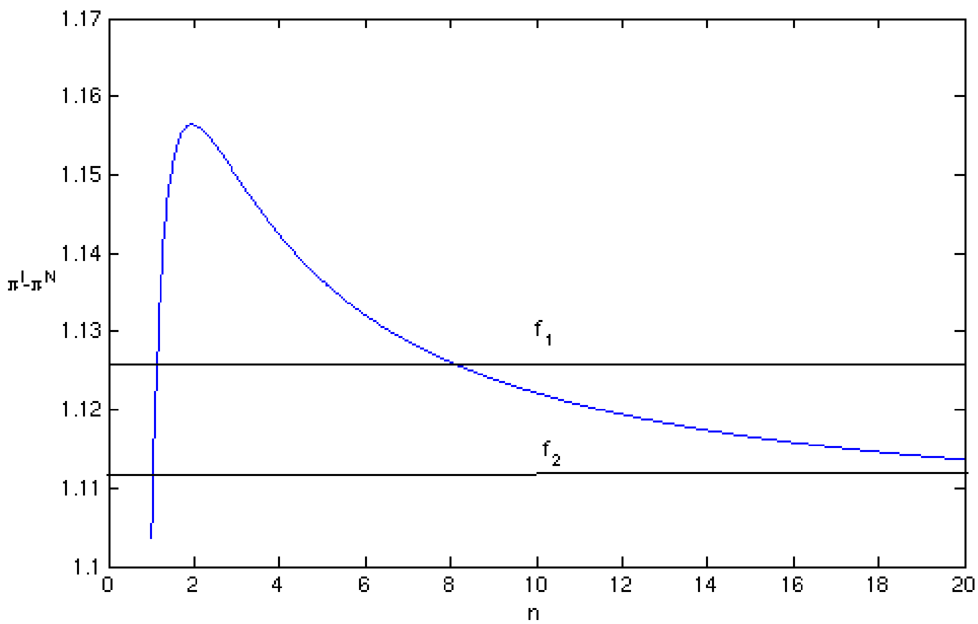

though, by modeling the upgrading cost as a fixed cost, I am unable to capture any differences in the intensity of innovation among the most competitive industries. If there were a variable component to the level and cost of upgrading, Figure 1 suggests that the intensity of upgrading efforts would be greatest for industries with intermediate n, in which $\pi^I(n) - \pi^N(n)$ is the greatest.

The inverse-U shape of the $\pi^I(n) - \pi^N(n)$ function results from two competing forces. The first is that as the number of competitors increases, the rents that the leader can earn fall as the leader's market share falls and the industry price is driven down closer to the leader's marginal cost.[2] When rents fall, so do the gains from upgrading, all else

[2]Note that if the leader upgrades, the price will converge to the followers' marginal cost for large n,

equal. This is the familiar Schumpeterian effect of competition on innovation, and it can be seen in the decreasing portion of $\pi^I - \pi^N$.

There is another effect at work, though, similar to what Aghion et al (2009) call the "escape from competition" effect. In a highly competitive industry, gaining a productive advantage over one's competitors can result in a large increase in market share. The gains from upgrading, then, are potentially large when n is high. This can be seen most starkly for the case $\gamma \geq \sigma/(\sigma-1)$. Recall that when γ is greater than the monopoly mark-up, a leader that upgrades can set the monopoly price and capture the entire market for the industry's product. A leader that does not upgrade, though, will have the same productivity as the followers. Since the Cournot equilibrium is symmetric, the leader's market share is equal to $(1/n)$. A monopolist, by already controlling the industry, stands to gain no market share by upgrading. On the other hand, a leader in a highly competitive industry with large n can go from having a miniscule market share to becoming a monopolist by upgrading. The "escape from competition" effect suggests that greater competition leads to greater innovative activity, and this can be seen in the upward-sloping portion of $\pi^I - \pi^N$. While $\gamma \geq \sigma/(\sigma-1)$ (so that the leader captures the entire market by upgrading) is an extreme case of the model, it is a useful one for illustrative purposes. The combination of the escape-from-competition effect and the Schumpeterian effect leads to the hump shape of $\pi^I - \pi^N$.[3]

4.3 Aggregation and equilibrium

Firm quantities are aggregated as follows. First, let $\pi_i(n)$ denote the variable profits earned by firm i operating in industry n, before the fixed cost of upgrading (if applicable) is deducted. Total variable industry profits $\pi(n)$ are given by:

$$\pi(n) = \sum_{i=1}^{n} \pi_i(n) \tag{4.5}$$

and not the leader's marginal cost. But it will be closer to the leader's marginal cost than it would have been for small n, though. See equation Section 2.3.

[3]Aghion et al (2009) use the phrase "escape from competition" in a slightly different sense from the way I use it. In their model, there are two producers, and the intensity of competition is related to how close the two producers' productivities are to each other. Upgrading helps one firm separate itself from the other. In my model, by contrast, the number of firms in the industry proxies for the industry's competitiveness.

Aggregate variable profits π then consists of the total amount of variable industry profits:

$$\pi = \int_1^\infty \pi(n) df(n) \tag{4.6}$$

With a labor supply equal to 1 and labor as the numeraire commodity, total labor income is equal to 1. Total income in the economy is the sum of labor income and aggregate profits, less upgrading costs. To account for the fixed costs that leaders pay to upgrade, I introduce an indicator function $x(\cdot) : [1, \infty) \mapsto \{0, 1\}$ that takes on value 1 if the leader in industry $n \geq 1$ decides to upgrade, and zero otherwise. The total fixed cost paid by all firms, F, is given by:[4]

$$F = \int_1^\infty x(n) f df(n)$$

The goods market clearing condition requires that total income be equal to total expenditures:

$$Pc = 1 + \pi - F \tag{4.7}$$

Let $\ell_i(n)$ denote the labor demanded by firm i in industry n. The total industry demand for labor, $\ell(n)$, is given by:

$$\ell(n) = \sum_{i=1}^{n} \ell_i(n)$$

Labor market clearing requires that the total labor demand be equal to the total labor supply:

$$\int_1^\infty \ell(n) df(n) = 1 \tag{4.8}$$

4.3.1 Definition of equilibrium

An equilibrium consists of prices $p(\cdot)$, quantities $q_i(\cdot)$, labor demands $\ell_i(\cdot)$, and an indicator function for upgrading decisions $x(\cdot)$ such that:

1. Given P, c, and $x(\cdot)$, $q_i(n)$ solves the profit-maximization problem of firm i in industry n, and $\ell_i(n)$ is the associated labor demand;

[4] I have not yet demonstrated that $x(\cdot)$ is integrable, but throughout this paper it will be.

2. For each industry n, $p(n)$ is given by the inverse demand function (1);

3. For each industry n, $x(n)$ maximizes the profits of the industry n leader:

$$\max_{x(n)\in\{0,1\}} x(n)(\pi^I(n) - f) + (1 - x(n))\pi^N(n)$$

4. The goods maret clearing condition (7) holds.

By Walras' Law, the labor market clearing condition is redundant.

The equilibrium of the model will not in general be unique. A leader's upgrading decision depends on aggregate quantities, which in turn depend on the upgrading decisions of leaders in other industries. Because upgrading involves paying a fixed cost to gain a discrete jump in productivity, many patterns of upgrading are possible across industries. Throughout this paper, I make the natural assumption that the equilibrium that holds is the one in which only the leaders with the strongest incentive to upgrade do so (if any firms upgrade at all). Formally, consider two industries n_1 and n_2. I consider equilibria in which for any n_1 and n_2, the following condition holds:

$$\pi^I(n_1) - \pi^N(n_1) > \pi^I(n_2) - \pi^N(n_2) \Rightarrow x(n_1) \geq x(n_2) \tag{4.9}$$

In other words, if the leader in industry n_1 has more to gain by upgrading than the leader in industry n_2, I will not consider equilibria in which the leader in n_2 upgrades but the leader in industry n_1 does not.

4.4 Social planner problem

To quantify the losses from allocative inefficiency in the benchmark model, we need to compare the equilibrium welfare to the welfare that arises from the social planner problem. I take real consumption, c, as the appropriate welfare measure. Denote by c^E the equilibrium welfare. Because of monopoly distortions, it will be lower than the social planner welfare, which I denote by c^{SP}.

The social planner must first make the upgrading decision $x(n)$ for the leader in each industry. Then, given the new distribution of productivities, the social planner chooses the welfare-maximizing output for each firm. The planner's problem can be solved in two steps. First, suppose that the planner has chosen to upgrade the technology of a fraction μ of all firms. Since all leading firms have the same initial productivity γz, the social planner will treat them as if they are identical. If the leader's technology is upgraded, then it will have a productivity advantage over the follower firms. And because the technology is linear, it will be the only active producer in the industry. If the leader's technology is not upgraded, then the pattern of production across firms within the industry is indeterminate. Since the total industry output will be known, we can assume that the leader produces the entire industry's output for the purposes of computing welfare. So whether or not a leader's technology is upgraded, we can treat each industry as if the leader is the only producer in the industry. Given μ, I compute the welfare-maximizing production level for each firm. Denote the maximized level of welfare by $c^{SP}(\mu)$. In the second step, the planner chooses the value μ^* that maximizes $c^{SP}(\mu)$. Welfare under the efficient social planner allocation is then given by $c^{SP} = c^{SP}(\mu^*)$.

The following assumption makes it easy to obtain an analytical solution to the social planner's problem but is not critical for the qualitative points being made. I will maintain it for the remainder of the paper.

Assumption 1. $\sigma = 2$.

For Step 1 of the social planner's problem, take μ as given. Given μ, the social planner chooses quantities that solve the following welfare maximization problem[5]

$$\text{Max} \left(\int_0^1 q(j)^{(\frac{\sigma-1}{\sigma})} dj \right)^{(\frac{\sigma}{\sigma-1})}$$

s.t.

$$\int_0^{\mu} \left(\frac{1}{\gamma^2}\right) q(j)dj + \int_{\mu}^{1} \left(\frac{1}{\gamma}\right) q(j)dj + \mu f \leq 1$$

The constraint is the resource constraint on labor, which is in unit supply. Solving the above problem is straightforward, and one can show that the resulting welfare is given by

$$c^{SP}(\mu) = \gamma\big(1 - \mu f\big)\big[\mu(\gamma - 1) + 1\big] \tag{4.10}$$

[5]Note that in this section I am reverting to the original notation in which industries are indexed by j. I do so because the leader is now the only active producer in each industry, so n does not matter.

The second step of the social planner's problem is to choose the welfare-maximizing level of μ:

$$\max_{\mu \in [0,1]} c^{SP}(\mu) = \gamma\big(1 - \mu f\big)\big[\mu(\gamma - 1) + 1\big]$$

The above problem will have an interior solution iff the following two conditions hold:

$$\text{(a) } \gamma > 1 + f, \text{ (b) } \gamma < 1 + \frac{f}{1 - 2f}$$

If (a) is violated, then $\mu^* = 0$. If (b) is violated, then $\mu^* = 1$. If neither is violated, then the solution is interior and given by

$$\mu^* = \frac{\gamma - 1 - f}{2f(\gamma - 1)} \tag{4.11}$$

What the above conditions say is that when the fixed cost, f, is small relative to the productivity gain γ, it is optimal for the planner to upgrade the technology of all firms. When the fixed cost is sufficiently large relative to γ, it is optimal for there to be no upgrading. For an intermediate range of f, it is optimal to have only some firms upgrade.

4.5 Allocative inefficiency

In the model, losses from allocative inefficiency result from the fact that firms' market power differs across industries. The central question of this section is: To what extent do firm innovation decisions amplify or attenuate distortions from firms having different market power? As I will show, upgrading amplifies losses from allocative efficiency. For example, when leaders in all but the least competitive industries upgrade (see Figure 2), upgrading increases the losses from allocative inefficiency because leaders in more competitive industries produce more efficiently and draw resources away from industries that are already producing too few goods relative to the social optimum.

I use the ratio of equilibrium welfare to social planner welfare to define a measure of allocative inefficiency, denoted by W_{loss}:

$$W_{loss} = 1 - \frac{c^E}{c^{SP}}$$

Table 2: Parameterization for Figure 2

Parameter	Value
σ	2
γ	1.01
z	1

For a particular set of parameters, I compute W_{loss} in the baseline model and compare it to the value of W_{loss} that arises in a version of the model in which firms do not have the opportunity to make investments in productivity growth. For the no-upgrading version of the model, W_{loss} is computed the same way as above, except now both c^E and c^{SP} are computed in a version of the model with no upgrading. If W_{loss} is higher in the no-upgrading model than in the baseline model, then upgrading mitigates monopoly distortions. If W_{loss} is higher in the no-upgrading model, then upgrading amplifies monopoly distortions.

The main exercise I will do will be to allow f to vary while holding all other

Figure 2: W_{loss} ratio as a function of f

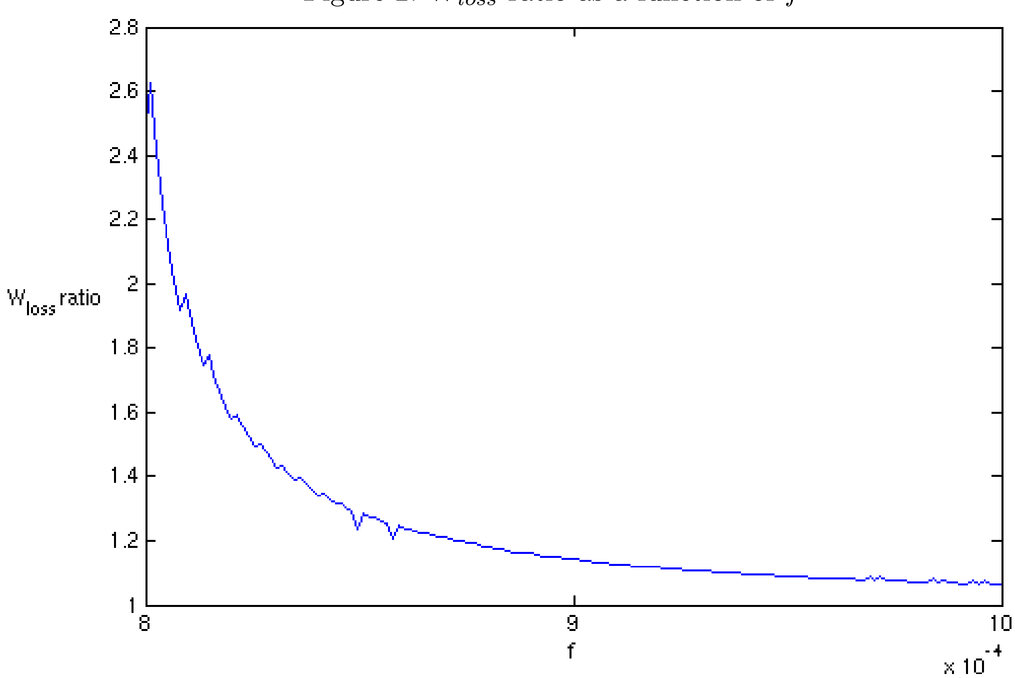

parameters fixed. The fixed parameters are given in Table 2. I have chosen values of f which ensure that only firms with an intermediate number of producers upgrade, thus mimicking the inverse-U relationship between competition and innovation. I assume that the number of competitors in an industry, n, is uniformly distributed on the interval [1,20].

Figure 2 shows the main results of the paper. The vertical axis measures the ratio of W_{loss} in the model with upgrading to W_{loss} in the version of the model in which there is no upgrading. (Note that the former depends on f while the latter does not, since upgrading is not allowed in the latter model.) What the figure shows is that the relative losses from allocative efficiency are larger when leaders are allowed to upgrade.

Moreover, the differences in W_{loss} between the two models are largest for low values of f. When f is low, many leaders upgrade, drawing many resources away from less-competitive, non-upgrading industries that are already producing below the social optimum. For values of f at the low range of the interval I consider, the relative losses from allocative inefficiency are over twice as large as in the no-upgrading version of the

Figure 3: Fraction of leaders that upgrade

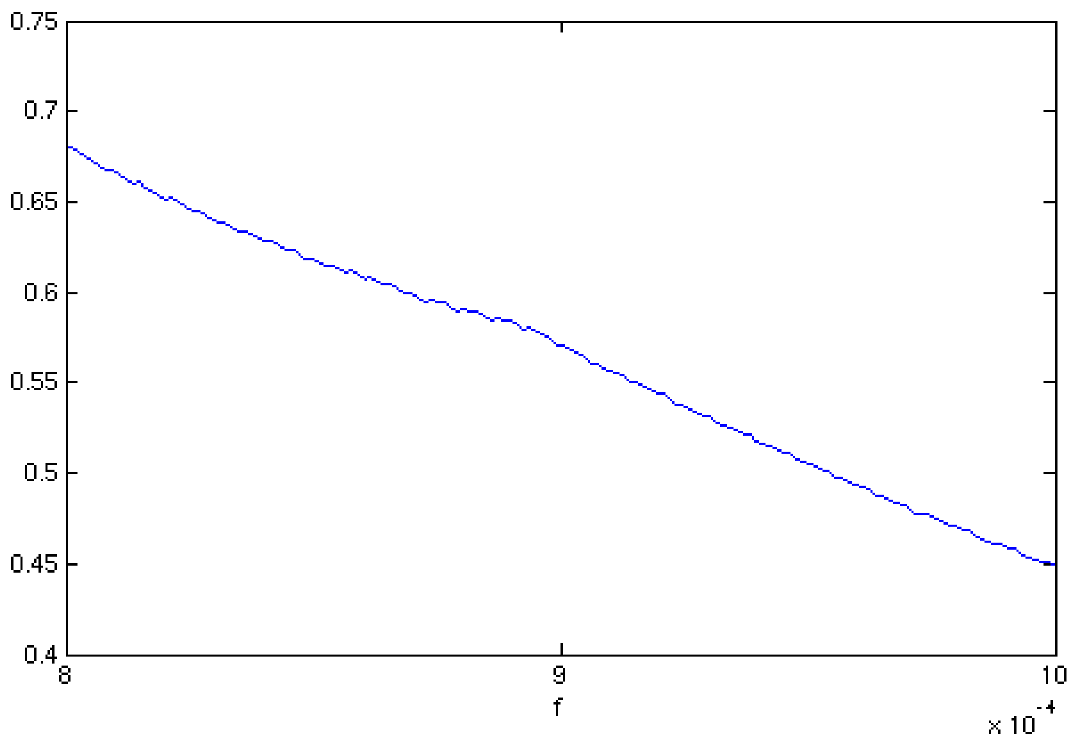

model. When f is high, fewer leaders upgrade, and the pull of resources away from uncompetitive industries is less strong. The relative losses from allocative inefficiency are then lower. Figure 3 shows how the fraction of leaders that upgrade declines as f increases.

4.6 Conclusion

This paper shows that technological upgrading decisions can increase distortions in environments where the level of competition differs across industries. The model I develop is well-suited to study the welfare implications because it captures the empirically observed inverse-U relationship between competition and innovation in a general equilibrium setting. In the numerical simulations I perform, differences in upgrading across firms amplify distortions arising from differences in market concerntration across industries by a factor of 2. Monopolists' low propensity to innovate is a well-known cost of monopoly in that they forego opportunities to produce at lower unit costs when it would be socially optimal for them to do so. My paper identifies an additional source of loss beyond monopolists' low productivity per se: Firms in more competitive industries draw resources away from low-competition industries, causing output in low-competition industries to fall further below the social optimum.

Bibliography

Abowd, J. M., Kramarz, F, & Margolis, D. N. (1999). High Wage Workers and High Wage Firms, *Econometrica*, 67, pp. 251333.

Aggarwal, R. K., & Samwick, A. A. (2006). Empire-builders and shirkers: Investment, firm performance, and managerial incentives. *Journal of Corporate Finance*, 12(3), 489-515.

Aghion, P., Blundell, R., Griffith, R., Howitt, P., & Prantl, S. (2009). The effects of entry on incumbent innovation and productivity. *The Review of Economics and Statistics*, 91(1), 20-32.

Arrow, K. (1962). Economic welfare and the allocation of resources for invention. In *The rate and direction of inventive activity: Economic and social factors* (pp. 609-626). NBER.

Atkeson, A., & Burstein, A. (2010). Innovation, Firm Dynamics, and International Trade. Federal Reserve Bank of Minneapolis Research Department Staff Report 444.

Bernard, A. B., Eaton, J., Jensen, J. B., & Kortum, S. (2003). Plants and Productivity in International Trade. *American Economic Review*, 1268-1290.

Bloom, N., Draca, M., & Van Reenen, J. (2011). Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity (No. w16717). NBER.

Bloom, N., & Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, 122(4), 1351-1408.

Bloom, N., Romer, P., Terry, S., & Van Reenen, J. (2013). A trapped factors model of innovation. CEP Discussion Paper, No. 1189.

Broda, C., & Weinstein, D. E. (2006). Globalization and the Gains from Variety. *The Quarterly Journal of Economics*, 121(2), 541-585.

Brown, C. & Medoff, J. (1989). The Employer Size-Wage Effect, *Journal of Political Economy*, 97, 102759.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms. *The American Economic Review*, 101(1), 304-340.

Caliendo, L. & Rossi-Hansberg, E. (2012). The Impact of Trade on Organization and Productivity, *The Quarterly Journal of Economics*, 127, 13931467.

Card, D. & DiNardo, J. E. (2002). Skill biased technological change and rising wage inequality: some problems and puzzles, Tech. rep., NBER.

Costantini, J., & Melitz, M. (2007). The dynamics of firm-level adjustment to trade liberalization. *The Organization of Firms in a Global Economy*, 107-141.

Davis, S. J., Haltiwanger, J., Katz, L.F., & Topel, R. (1991). Wage Dispersion between and within U.S. Manufacturing Plants, 1963-86, *Brookings Papers on Economic Activity*. Microeconomics, 1991, pp. 115200.

Doms, M., Dunne, T., & Troske, K.R. (1997). Workers, wages, and technology, *The Quarterly Journal of Economics*, 112, 253290.

Dunne, T., Foster, L., Haltiwanger, J., & Troske, K. (2002). Wage and Productivity Dispersion in U.S. Manufacturing: The Role of Computer Investment, IZA Discussion Papers 563, Institute for the Study of Labor (IZA). 18

Dunne, T. & Schmitz, J. A (1995). Wages, Employment Structure and Employer Size-Wage Premia: Their Relationship to Advanced-Technology Usage at US Manufacturing Establishments, *Economica*, 62, 89107.

Dyck, A., & Zingales, L. (2004). Private benefits of control: An international comparison. *The Journal of Finance*, 59(2), 537-600.

Edwards, S., & Lederman, D. (1998). The political economy of unilateral trade liberalization: the case of Chile (No. w6510). NBER.

Fernandes, A. M. (2007). Trade policy, trade volumes and plant-level productivity in Colombian manufacturing industries. *Journal of International Economics*, 71(1), 52-71.

Garicano, L. (2000). Hierarchies and the Organization of Knowledge in Production, *Journal of Political Economy*, 108, 874904.

Grobovsek, J. (2013). Management and Aggregate Productivity, Mimeo, University of Edinburgh.

Harberger, A. C. (1954). Monopoly and resource allocation. *The American Economic Review*, 44(2), 77-87.

Harris, M., & Raviv, A. (1979). Optimal incentive contracts with imperfect information. *Journal of Economic Theory*, 20(2), 231-259.

Haskel, J. (1998). The Skilled Wage Premium and the Growth of Small Firms in UK Manufacturing, *LABOUR*, 12, 221238.

Helpman, E., Itskhoki, O., & Redding, S. (2010). Inequality and Unemployment in a Global Economy, *Econometrica*, 78, 12391283.

Holmes, T. J., & Schmitz, J. A. (2001). A gain from trade: From unproductive to productive entrepreneurship. *Journal of Monetary Economics*, 47(2), 417-446.

Holmes, T. J., & Schmitz, J. A. (2010). Competition and productivity: a review of evidence (Vol. 439). DIANE Publishing.

Holmes, T. J., Levine, D. K., & Schmitz, J. A. (2012). Monopoly and the Incentive to Innovate When Adoption Involves Switchover Disruptions. *American Economic Journal: Microeconomics*, 4(3), 1-33.

Holmstrom, B. (1979). Moral hazard and observability. *The Bell Journal of Economics*, 74-91.

Idson, T. L. & Oi, W. Y. (1999). Workers Are More Productive in Large Firms, *American Economic Review*, 89, 104108.

Katz, L. F. & Summers, L. H. (1989). Can Interindustry Wage Differentials Justify Strategic Trade Policy? in *Trade Policies for International Competitiveness*, NBER, NBER Chapters, 85124.

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

Liu, L. (1993). Entry-exit, learning, and productivity change Evidence from Chile. *Journal of Development Economics*, 42(2), 217-242.

Liu, L., Roberts, M., & Tybout, J. (1996). Productivity Growth in Colombia and Chile: Panel-based Evidence on the role of Entry, Exit and Learning. *Industrial Evolution in*

Developing Countries: Micro Patterns of Turnover, Productivity, and Market Structure.

Muendler, M. A. (2004). Trade, technology, and productivity: A study of Brazilian manufacturers, 1986-1998.

Pavcnik, N. (2002). Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants. *Review of Economic Studies*, 69(1), 245-276.

Pavcnik, N. (2003). What explains skill upgrading in less developed countries? *Journal of Development Economics*, 71, 311328.

Raith, M. (2003). Competition, risk and managerial incentives. *American Economic Review*, 93, 1425-1436.

Reksulak, M., Shughart, W. F., & Tollison, R. D. (2008). Innovation and the opportunity cost of monopoly. *Managerial and Decision Economics*, 29(8), 619-627.

Sappington, D. (1983). Limited liability contracts between principal and agent. *Journal of Economic Theory*, 29(1), 1-21.

Schmidt, K. M. (1997). Managerial incentives and product market competition. *The Review of Economic Studies*, 64(2), 191-213.

Schmitz, J. A. (2005). What determines productivity? Lessons from the dramatic recovery of the US and Canadian iron ore industries following their early 1980s crisis. *Journal of Political Economy*, 113(3), 582-625.

Schumpeter, J. (1942). *Capitalism, Socialism, and Democracy.* New York: Harper and Row.

Sidak, J. G., & Teece, D. J. (2009). Dynamic competition in antitrust law. *Journal of Competition Law and Economics*, 5(4), 581-631.

Sivadasan, J. (2009). Barriers to competition and productivity: Evidence from India. *The BE Journal of Economic Analysis & Policy*, 9(1), 42.

Syverson, C. (2004). Product substitutability and productivity dispersion. *Review of Economics and Statistics*, 86(2), 534-550.

Yeaple, S. R. (2005). A simple model of firm heterogeneity, international trade, and wages, *Journal of International Economics*, 65, 120.

Zweimuller, J. & Winter-Ebmer, R. (2003). On-the-job-training, job search and job mobility, *Swiss Journal of Economics and Statistics*, 139, 563576.

Appendix A

Appendix to Chapter 2

Proof of Proposition 1.

Given s_H and s_L, the manager's problem is

$$\rho^*(s_H, s_L) = \max_{\rho \in [0,1]} \rho s_H \pi_H + (1-\rho) s_L \pi_L - A z^{\sigma-1} \exp(b\rho) + \lambda z^{\sigma-1}$$

Throughout I will consider only interior solutions to the manager's problem. For appropriate choices of A and b the solution will be strictly between 0 and 1. The manager's problem is strictly concave in ρ, so the optimal choice of ρ, $\rho^*(s_H, s_L)$, solves the first-order condition from the manager's problem:

$$s_H \pi_H - s_L \pi_L - A b z^{\sigma-1} \exp(b\rho) = 0 \tag{A.1}$$

Rearranging this gives

$$\rho^*(s_H, s_L) = \Big(\frac{1}{b}\Big) \log\Big[\Big(\frac{1}{Ab}\Big) z^{1-\sigma} (s_H \pi_H - s_L \pi_L)\Big] \tag{A.2}$$

The owners' problem can now be written as

$$\max_{s_h, s_L} \rho^*(s_H, s_L)[(1-s_H)\pi_H] + (1-\rho^*(s_h, s_L))[(1-s_L)\pi_L]$$

$$\text{s.t.}$$

$$s_H \geq 0$$

$$s_L \geq 0$$

$$s_H \leq 1$$

$$s_L \leq 1$$

$$\rho^*(s_H, s_L)s_H\pi_H + (1-\rho^*(s_H, s_L))s_L\pi_L - e(\rho^*, z) + \lambda z^{\sigma-1} \geq w_h$$

$$\rho^*(s_H, s_L) = \Big(\frac{1}{b}\Big)\log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s_H\pi_H - s_L\pi_L)\Big]$$

To prove the first part of Proposition 1, that $s_L = 0$, I need the following lemma.

Lemma 6. *Under the optimal contract,* $(1-s_H)\pi_H \geq (1-s_L)\pi_L$.

Proof. Let $s = (s_H, s_L)$ be an optimal contract. Suppose, to show a contradiction, that $(1-s_H)\pi_H < (1-s_L)\pi_L$. Now consider an alternate contract $\bar{s} = (s_H - \epsilon, s_L)$, where ϵ is small. To show the contradiction, it suffices to show that owners have higher expected income under the contract $\bar{s}$. Denote by $W(s)$ the expected income that owners receive under contract s. Subtracting the owners' expected income under $\bar{s}$ from the owners' income under s gives

$$\begin{aligned} W(s) - W(\bar{s}) &= \rho^*(s)(1-s_H)\pi_H + (1-\rho^*(s))(1-s_L)\pi_L - \rho^*(\bar{s})(1-s_H+\epsilon)\pi_H \\ &\quad + (1-\rho^*(\bar{s}))(1-s_L)\pi_L \\ &= (\rho^*(s) - \rho^*(\bar{s}))\Big[(1-s_H)\pi_H - (1-s_L)\pi_L\Big] - \rho^*(\bar{s})\epsilon\pi_H \\ &< 0. \end{aligned}$$

It is immediate from the expression for ρ^* in (23) that $\rho^*(s) - \rho^*(\bar{s}) > 0$. And by assumption, $(1-s_H)\pi_H - (1-s_L)\pi_L < 0$. The inequality in the third line then follows. Since $W(s) - W(\bar{s}) < 0$, s cannot be an optimal contract, so there is a contradiction and the lemma is proved.

Now I am in a position to prove that the manager's limited-liability constraint $s_L \geq 0$ binds in the low-profit state. The proof is by contradiction. Suppose that $s = (s_H, s_L)$ is an optimal contract and that $s_L > 0$. To show that this cannot be an optimal contract, I show that the contract $\bar{s} = (s_H, s_L - \epsilon)$ offers higher expected value to the owners, where ϵ is small. The difference in the owners' expected value between the two contracts is

$$W(s) - W(\bar{s}) = \rho^*(s)(1-s_H)\pi_H + (1-\rho^*(s))(1-s_L)\pi_L - \rho^*(\bar{s})(1-s_H)\pi_H$$
$$+ (1-\rho^*(\bar{s}))(1-s_L+\epsilon)\pi_L$$
$$= (\rho^*(s) - \rho^*(\bar{s}))\Big[(1-s_H)\pi_H - (1-s_L)\pi_L\Big] - \rho^*(\bar{s})\epsilon\pi_L$$
$$< 0.$$

From the expression for ρ^* in (23), we see that $(\rho^*(s) - \rho^*(\bar{s}))$. And from Lemma 1, we have that $(1-s_H)\pi_H \geq (1-s_L)\pi_L$. The inequality in line 3 then follows. With $W(s) - W(\bar{s}) < 0$, the contradiction is shown, and therefore the optimal contract must have $s_L = 0$.

To establish part (2) of Proposition 1, consider once more the owners' maximization problem. The first-order Kuhn-Tucker condition with respect to s_H is

$$\frac{1}{b}\Big(\frac{\pi_H}{s_H\pi_H - s_L\pi_L}\Big)[(1-s_H)\pi_H - (1-s_L)\pi_L] - \frac{1}{b}\log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s_H\pi_H - s_L\pi_L)\Big]\pi_H$$
$$+\mu_M(H) - \mu_O(H) = 0$$

where $\mu_M(H)$ is the multiplier on the limited-liability constraint $s_H \geq 0$ and $\mu_O(H)$ is the multiplier on the owners' limited-liability constraint $s_H \leq 1$. The fact that the owners' limited-liability constraint does not bind is established in the following lemma.

Lemma 7. *Under the optimal contract, the limited-liability constraint on owners does not bind in the high-profit state. That is, $s_H < 1$.*

Proof. It has already been established that an optimal contract has $s_L = 0$. The proof follows from the fact that the contract $\bar{s} = (1-\epsilon, 0)$ offers the owners greater expected value than the contract $s = (1, 0)$. Showing this is straightforward.

Having established that $s_H < 1$ under the optimal contract, there are two cases left to consider. The first is when the manager's limited-liability constraint on s_H does not bind. In this case, the optimal s_H solves equation (26), with $\mu_M(H)$ and $\mu_O(H)$ set equal to zero. (This follows from the complementary slackness conditions.) The solution is the value of s^* given in part (2) of Proposition 1. If $s^* < 0$, then the limited-liability constraint binds and $s_H = 0$. Hence, under the optimal contract, $s_H = \max\{s^*, 0\}$.

Part (3) of Proposition (2) follows from subsituting $s_H = \max\{s^*, 0\}$ and $s_L = 0$ into (23). $\square$

Proof of Proposition 2.

For a struggling firm, the manager's problem is

$$\rho^*(s_H) = \max_{\rho \in [0,1]} \rho s_H \pi_H + \rho \lambda z^{\sigma-1} - A z^{\sigma-1} \exp(b\rho)$$

The maximization problem is concave in ρ, so (ignoring corner solutions) the manager's optimal choice of ρ will satisfy the first-order condition

$$s_H \pi_H + \lambda z^{\sigma-1} - Abz^{\sigma-1} \exp(b\rho) = 0 \tag{A.3}$$

Rearranging this gives the manager's optimal choice of ρ:

$$\rho^*(s_H) = \Big(\frac{1}{b}\Big) \log\Big[\Big(\frac{1}{Ab}\Big) z^{1-\sigma}(s_H \pi_H + \lambda z^{\sigma-1})\Big] \tag{A.4}$$

With the above expression for ρ^* in hand, the owners' problem can be written as

$$\max_{s_H} \rho^*(s_H)[(1 - s_H)\pi_H]$$

$$\text{s.t.}$$

$$s_H \geq 0$$

$$s_H \leq 1$$

$$\rho^*(s_H)[s_H \pi_H + \lambda z^{\sigma-1}] - A z^{\sigma-1} \exp(b\rho) \geq w_h$$

$$\rho^*(s_H) = \Big(\frac{1}{b}\Big) \log\Big[\Big(\frac{1}{Ab}\Big) z^{1-\sigma}(s_H \pi_H + \lambda z^{\sigma-1})\Big]$$

By assumption, the participation constraint does not bind, so we can ignore it. Also, by an argument parallel to the one in the proof to Proposition 1, one can show that the limited-liability constraint on owners does not bind. We can therefore ignore this constraint going forward as well. The first-order Kuhn-Tucker condition for the owners' problem is

$$\frac{1}{b}\Big(\frac{\pi_H}{s_H\pi_H+\lambda z^{\sigma-1}}\Big)(1-s_H)\pi_H-\frac{1}{b}\log\Big[\Big(\frac{1}{Ab}\Big)z^{1-\sigma}(s_H\pi_H+\lambda z^{\sigma-1})\Big]+\mu_M(H)=0 \quad \text{(A.5)}$$

where $\mu_M(H)$ is the multiplier on the manager's limited-liability constraint in the high-profit state. If the limited-liability constraint does not bind, then by the complementary slackness condition the optimal choice of s_H solves equation (29) with $\mu_M(H)$ set to 0. If the limited-liability constraint does bind, then it is immediate that $s_H = 0$. This completes the proof of part (1) of Proposition 2. Part (2) comes from substituting the value of s_H found in part (1) into equation (28). □

The case where the participation constraint binds

In this section, I characterize the optimal contract between a manager and firm owners when the participation constraint binds. Struggling firms are the easiest to analyze and will be considered first. I then turn to thriving firms.

Struggling firms

Since profits in the low-productivity state are zero, the contract is a one-dimensional object which only specifies the share of profits that a manager receives in the high state. When the participation constraint binds, we have

$$\rho^*(s_H)(s_H\pi_H+\lambda z^{\sigma-1})-Az^{\sigma-1}\exp(b\rho^*(s_H))=w_h \quad \text{(A.6)}$$

The left-hand side of this equation is strictly increasing in s_H, so there is a unique value of s_H that satisfies it. If the above equality holds for some $s_H > 1$, then no contract can be formed (the owners' limited-liability constraint is violated) and the firm exits. Note that when the participation constraint binds, the manager receives a larger share of profits than she otherwise would if the participation constraint did not bind. As a result, the manager exerts more effort when the participation constraint binds. This is another mechanism through which increased competition can raise productivity in the model. Lower tariffs erode profits for domestic firms, causing the participation constraint to bind for some of them: The share of profits that the manager had previously been receiving is no longer sufficient to satisfy the participation constraint. As a result, the managers receive a larger share of profits and work harder to raise productivity.

Thriving firms

The following proposition summarizes the contract received by the manager of a thriving firm whose participation constraint binds.

Proposition 6. *Suppose a firm with productivity z is thriving and that the manager's participation constraint binds. Let $\bar{s}$ be defined so that $(1-\bar{s})\pi_H = \pi_L$. Also, let s^* be the solution to*

$$\rho^*(s^*, 0)s^*\pi_H - e(z, \rho^*(s^*, 0)) + \lambda z^{\sigma-1} = w_h$$

The manager's contract is then summarized as follows:

(1) If $0 \leq s^* < \bar{s}$, then the owners choose the contract $(s_H, s_L) = (s^*, 0)$.

(2) If $1 \geq s^* \geq \bar{s}$, then the owners choose the unique contract (s_H, s_L) such that:

$$(1 - s_H)\pi_H = (1 - s_L)\pi_L$$

$$\rho^*(s_H, s_L)s_H\pi_H - (1 - \rho(s_H, s_L))s_L\pi_L - e(z, \rho^*(s_H, s_L)) + \lambda z^{\sigma-1} = w_h$$

A formal proof is available upon request. Here I just provide an intuitive explanation. To satisfy the manager's participation constraint, the owners must give payments to managers larger than those specified in Proposition 1. The owners will first pay the manager a larger share of the profits when productivity is high than they would if they were not constrained, and keep $s_L = 0$. The reason is that, as long as s_H is not too high, the owners receive more income in the high-profit state than in the low-profit state. They therefore want to give the manager strong incentives to exert effort. Doing so entails punishing the manager as much as possible in the low-profit state by setting $s_L = 0$. The value of s_H (given $s_L = 0$) that satisfies the manager's particpation constraint with equality is given by s^*.

However, if the value of s_H required to satisfy the participation constraint (given $s_L = 0$) is too high, then manager extracts all of the surplus from the productivity gains, and then some. When $(1-s_H)\pi_H < \pi_L$, the owners would prefer that the manager exert no effort at all. Once all of the surplus from the productivity gains has been extracted by the manager, the owners choose the contract that sets $(1 - s_H)\pi_H = (1 - s_L)\pi_L$

and satisfies the manager's participation constraint with equality. Any other contract would mean that the manager's payments are highest in the state in which the owners' payments are the lowest. (If $s_H\pi_H > s_L\pi_L$, then $(1 - s_H)\pi_H < (1 - s_L)\pi_L$, and vice versa.) As a result, the incentive structure will encourage the manager to try to attain the state that the owners least prefer, which cannot be optimal.

Table A1: Industry Trade-Orientation Classifications

Trade Orientation	ISIC code	Industry Name
Import-competing	321	Textiles
	351	Industrial chemicals
	355	Rubber products
	361	Pottery, china, earthenware
	362	Glass products
	381	Fabricated metal products
	382	Machinery, excl. electrical
	383	Machinery, electric
	384	Transport equipment
	385	Professional and scientific equipment
	390	Other manufactured products
Non-traded	311	Food products
	313	Beverages
	314	Tobacco
	322	Wearing apparel, excl. footwear
	323	Leather products
	324	Footwear, excl. rubber or plastic
	332	Furniture, excl. metal
	342	Printing and publishing
	352	Other chemicals
	353	Petroleum refineries
	354	Misc. petroleum and coal products
	356	Plastic products
	369	Other non-metallic mineral products
	371	Iron and steel
Export-oriented	331	Wood products, excl. furniture
	341	Paper and products
	372	Non-ferrous metals

Appendix B

Appendix to Chapter 3

B.1 Select Proofs

B.1.1 Proof of Lemma 1

We will first prove by induction on x that each manager at level $x > 0$ is able to solve $(R\lambda)^x(1-\lambda)$ tasks. Consider first the measure of tasks that workers solve: by assumption each worker receives a measure one of tasks of which he is unable to solve a fraction λ which he passes on to managers in level $x = 1$. Now, each manager at level $x = 1$ receives R direct reports. Therefore each manager at level $x = 1$ receives a measure $R\lambda$ of tasks, out of which they solve a fraction $1-\lambda$ and passes on the remaining to managers at level $x = 2$.

Consider a manager at level x and assume that he solves a measure $(R\lambda)^x(1-\lambda)$ of tasks. A manager at level $x+1$ receives R direct reports, each of which passes on a measure $\lambda(R\lambda)^x$ of unsolved tasks. Of this measure he is able to solve a fraction $1-\lambda$. This implies that $(R\lambda)^{x+1}(1-\lambda)$.

Since there are R^{X-x} managers in level x, the total number of tasks $\mathcal{T}_X$ is given by

$$\begin{aligned}\mathcal{T}_X &= \sum_{x=0}^{X} R^{X-x}(R\lambda)^x(1-\lambda) \\ &= R^X(1-\lambda^{X+1}).\end{aligned}$$

B.1.2 Proof of Lemma 2

First we will prove that $SP_{X+1,R} > SP_{X,R}$. For this it is sufficient to prove that

$$\frac{R^{X+1}}{R^{X+1}-1}(1-\lambda^{X+1}) > \frac{R^X}{R^X-1}(1-\lambda^X),$$

which is equivalent to proving that

$$R\frac{\sum_{x=0}^{X}\lambda^x}{\sum_{x=0}^{X-1}\lambda^x} > \frac{\sum_{x=0}^{X}R^x}{\sum_{x=0}^{X-1}R^x}. \tag{B.1}$$

(B.1) holds if and only if

$$\begin{aligned} R\left(\sum_{x=0}^{X-1}R^x\right)\left(\sum_{x=0}^{X}\lambda^x\right) &> \left(\sum_{x=0}^{X}R^x\right)\left(\sum_{x=0}^{X-1}\lambda^x\right) \\ \iff \lambda^X\sum_{x=1}^{X-1}R^x &> \sum_{x=0}^{X-1}\lambda^x \\ \iff \sum_{x=0}^{X-1}\lambda^x\left(R\lambda\right)^{X-x} &> \sum_{x=0}^{X-1}\lambda^x. \end{aligned} \tag{B.2}$$

Since $R\lambda > 1$, λ^x on the left hand side of (B.2) is multiplied by a term greater than one for all $x \in \{0, \ldots, X-1\}$. Therefore the left hand side is strictly greater than the right hand side and the proof follows.

Now we prove that $SP_{X,R+1} > SP_{X,R}$. For this it is sufficient to prove that

$$R\frac{(R+1)^X}{(R+1)^X-1} > (R-1)\frac{R^X}{R^X-1},$$

which is equivalent to proving that

$$R(R^X-1)(R+1)^X > (R-1)R^X((R+1)^X-1). \tag{B.3}$$

(B.3) holds if and only if

$$(R^X-R)(R+1)^X + (R^{X+1}-R^X) > 0,$$

which holds if $R > 1$. Notice that $R\lambda > 1$ and $\lambda \in (0,1)$ imply $R > 1$, which completes the proof.

B.1.3 Proof of Proposition 3

We will prove the result by first considering levels of width Δ and calculating by induction the total number of tasks solved in $AC(x)$ as $\Delta \to 0$. In each level $x' \in [0,x]$ of area of command $AC(x)$ there are $R^{x-x'}$ employees. Workers in $AC(x,n)$ receive a measure 1 of tasks and are able to solve a fraction $\Delta(1-\lambda)$ of them. Each manager in level Δ receives $R^{\Delta}(1-\Delta(1-\lambda))$ tasks and is able to solve a fraction $\Delta(1-\lambda)$. Since there is a measure $R^{x-\Delta}$ of managers at level Δ, the total measure of tasks that managers at level Δ solve is $R^x(1-\Delta(1-\lambda))\Delta(1-\lambda)$. Each employee in level 2Δ manages a width of R^{Δ} direct reports, each of whom passes on $R^{\Delta}\left(1-\Delta(1-\lambda)\right)^2$ unsolved tasks. Each manager in level 2Δ is able to solve a fraction $\Delta(1-\lambda)$ of those tasks. There is a measure $R^{x-2\Delta}$ of managers at level 2Δ, so the total measure of tasks that these managers solve is $R^x(1-\Delta(1-\lambda))^2\Delta(1-\lambda)$.

Now assume that managers at level $i\Delta$ solve a measure $R^{i\Delta}(1-\Delta(1-\lambda))^i\Delta(1-\lambda)$ of tasks. There are $R^{x-i\Delta}$ managers at this level. So the total measure of tasks solved by managers at this level is $R^x(1-\Delta(1-\lambda))^i\Delta(1-\lambda)$. Each employee at level $(i+1)\Delta$ manages a width of R^{Δ} direct reports, each of whom passes on $R^{i\Delta}\left(1-\Delta(1-\lambda)\right)^{i+1}$ unsolved tasks. Each manager in level $(i+1)\Delta$ is able to solve a fraction $\Delta(1-\lambda)$ of those tasks. Since there are $R^{x-(i+1)\Delta}$ managers at level $i+1$, the total measure of tasks solved by these managers is $R^x(1-\Delta(1-\lambda))^{i+1}\Delta(1-\lambda)$. Therefore the total measure of tasks solved in area of command $AC(x)$ is

$$
\begin{aligned}
\mathcal{T}(x) &\equiv \lim_{\Delta\to 0}(1-\lambda)R^x\Delta\sum_{i=0}^{\left\lfloor \frac{x}{\Delta}\right\rfloor}\left(1-\Delta(1-\lambda)\right)^i \\
&= \lim_{\Delta\to 0} R^x(1-(1-\Delta(1-\lambda))^{\left\lfloor \frac{x}{\Delta}\right\rfloor+1}) \\
&= R^x\left(1-e^{-x(1-\lambda)}\right),
\end{aligned}
$$

where the last line is due to the fact that

$$\lim_{\Delta \to 0} \frac{x}{\Delta} \ln\left(1 - \Delta(1-\lambda)\right) = -x(1-\lambda).$$

B.1.4 Proof of Proposition 4

Recall from Proposition 3 that the number of tasks a firm solves is given by $\mathcal{T}(X) = R^X\left(1 - e^{-X(1-\lambda)}\right)$. Equilibrium condition 2 requires that consumers be indifferent between obtaining any education levels which are employed with positive measure. This implies that $w(x) = w_E h(x) + w$. Now, for an educator it must hold that $w_E = w_E h_E + w$, so $w_E = \frac{w}{1-h_E}$. Therefore

$$w(x) = \frac{w}{1-h_E} h(x) + w. \tag{B.4}$$

Plugging (B.4) and the labor requirements at each level into the firm's problem we get:

$$\begin{aligned}
&\max_X Y(\mathcal{T}(X), z) - \int_x R^{X-x}\left(\frac{w}{1-h_E}h(x) + w\right)dx \\
=&\max_X Y(\mathcal{T}(X), z) - R^X \int_x we^{-x(1-\lambda)}dx \\
=&\max_X Y(\mathcal{T}(X), z) - w\mathcal{T}(X),
\end{aligned} \tag{B.5}$$

where the second line is a consequence of Assumption 1.

Assumption 2 on the span of control implies that $\mathcal{T}(X)$ is strictly increasing in X.[1] Therefore (B.5) is equivalent to

$$\max_{\mathcal{T}} Y(\mathcal{T}, z) - w\mathcal{T}.$$

B.1.5 Proof of Proposition 5

Proposition 4 implies that the solution to (3.5) is sufficient to derive a function $X^* : \mathcal{Z} \to \mathbb{R}_+$ that satisfies the equilibrium conditions given wages.

[1] $\mathcal{T}'(X) = R^X\left(\ln R - e^{-X(1-\lambda)}\left(\ln R - (1-\lambda)\right)\right)$. Since $\ln R > 1-\lambda$, $1-\lambda > 0$ and $e^{-X(1-\lambda)} < 1$, then $\mathcal{T}'(X) > 0$.

Denote by $\mathcal{T}^*(z,w)$ the solution to (3.5). The regularity conditions on $Y(\mathcal{T},w)$ imply that $\mathcal{T}^*(z,w)$ is strictly decreasing in w, $\lim_{w\to\infty}\mathcal{T}^*(z,w)=0$ and $\lim_{w\to 0}\mathcal{T}^*(z,w)=\infty$ since $\mathcal{T}^*(z,w)$ satisfies

$$Y_{\mathcal{T}}(\mathcal{T}^*(z,w),z)=w.$$

Denote by $X^*(z,w)$ the value of X that achieves $\mathcal{T}^*(z,w)$. Notice that $X^*(z,w)$ is also strictly decreasing in w. Let

$$\begin{aligned} H(w) &\equiv \int_z\int_x l(x,X^*(z,w))\left(1+\frac{h(x)}{1-h_E}\right)f(z)dxdz \\ &= \int_z\int_x R^{X^*(z,w)-x}R^x e^{-x(1-\lambda)}f(z)dxdz \\ &= \frac{1}{1-\lambda}\int_z \mathcal{T}^*(z,w)f(z)dz. \end{aligned}$$

Then the regularity conditions on $Y(\mathcal{T},w)$ imply that $H(w)$ is strictly decreasing in w, $\lim_{w\to\infty}H(w)=0$ and $\lim_{w\to 0}H(w)=\infty$. Since the labor market clearing condition implies $L=H(w)$, the result follows.

B.1.6 Proof of Lemma 4

Plugging in $h(x)$ from Assumption 1 and the definition of $l(x,X)$ we get

$$SP(X)=\ln R\frac{1}{1-\lambda}\frac{R^X}{R^X-1}(1-e^{-X(1-\lambda)}).$$

Consider

$$g(X)\equiv\frac{R^X}{R^X-1}(1-e^{-X(1-\lambda)}).$$

It is sufficient to prove that $g'(X)>0$. Now

$$g'(X)=\frac{R^X}{e^{(1-\lambda)X}(R^X-1)^2}\left((1-\lambda)(R^X-1)-\ln R\left(e^{X(1-\lambda)}-1\right)\right).$$

It is sufficient to prove that $\widehat{g}(X)\equiv(1-\lambda)(R^X-1)-\ln R\left(e^{X(1-\lambda)}-1\right)>0$ for all $X>0$. First notice that $\widehat{g}(0)=0$. Now,

$$\widehat{g}'(X)=(1-\lambda)\ln R\left(R^X-e^{X(1-\lambda)}\right),$$

so $\ln R > 1-\lambda$ implies that $\widehat{g}'(X) > 0$.

B.1.7 Proof of Lemma 5

Plugging in $h(x)$ from Assumption 1 and the definition of $l(x, X)$ we get

$$SP(X) = \ln R \frac{1}{1-\lambda} \frac{R^X}{R^X - 1}(1 - e^{-X(1-\lambda)}).$$

Let $g(R) = \ln R \frac{R^X}{R^X-1}$. It is sufficient to prove that $g'(R) > 0$. Now

$$g'(R) = \frac{R^{X-1}}{(R^X-1)^2}\left(R^X - 1 - \ln R^X\right).$$

Since $x - 1 > \ln x$ for any $x > 0$, the result follows.