

Monotone Comparative Statics for the Industry Composition

Peter Arendorf Bache Anders Laugesen*

Working paper: June 17, 2015

Abstract

We let heterogeneous firms face decisions on a number of complementary activities in a monopolistically-competitive industry. The endogenous level of competition and selection regarding entry and exit of firms introduces a wedge between monotone comparative statics (MCS) at the firm level and MCS for the industry composition. The latter phenomenon is defined as first-order stochastic dominance shifts in the equilibrium distributions of all activities across active firms. We provide sufficient conditions for MCS at both levels of analysis and show that we may have either type of MCS without the other. It is therefore possible that firm-level complementarities manifest themselves more clearly at the industry level than at the firm level during comparative statics. This turns out to be the case for a large number of models and shocks considered in the recent trade literature for which we provide strong, novel, and testable predictions.

Keywords: Complementary Activities; Firm Heterogeneity; Monotone Comparative Statics; Markups; Firm-Size Distribution; International Trade

JEL Classifications: D21; F12; F61; L11; L22

*Bache works at the Ministry of Finance of Denmark. Laugesen (corresponding author) works at the Department of Economics and Business, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark. Email: alaugesen@econ.au.dk. We have greatly benefited from discussions with Marc Melitz, J. Peter Neary, and Allan Sørensen. We also thank Martin Alfaro, Torben M. Andersen, Pol Antràs, Kristian Behrens, Donald Davis, Hartmut Egger, John Feddersen, Alexander Koch, John Morrow, Mathieu Parenti, Horst Raff, John Roberts, Bo Sandemann, Philipp Schröder, Siri Strandenes, Jacques Thisse, and participants at numerous seminars and conferences.

1 Introduction

More than two decades ago, Milgrom and Roberts (1990a) argued that strategy and structure in modern manufacturing firms reflect widespread complementarities among many diverse activities undertaken by firms. Drawing on Topkis (1978), they emphasised how such complementarities make firms' decisions exhibit monotone comparative statics (MCS). That is, the optimal levels of the activities are monotonic in parameters of the profit maximisation problem that influence the set of available activities or the attractiveness of these activities, all else equal. Since this seminal contribution, the monotonicity theorems developed by Topkis (1978), Milgrom and Shannon (1994), and Athey (2002) have been central in comparative statics of firms. One reason is their virtue of focusing on the properties of the optimisation problem that are essential for obtaining MCS and doing away with superfluous assumptions. For instance, these monotonicity theorems allow activities to be discrete choice variables and the profit function to be nonconcave, nondifferentiable, and discontinuous. When applying these monotonicity theorems, one typically assumes that the competitive environment is exogenous.¹ While being a natural starting point, such an analysis is not entirely satisfactory when studying exogenous shocks that affect all firms in an industry. In this case, firms are not only directly affected by the exogenous changes but also indirectly affected through changes in the competitive environment.

In this paper, we apply the monotonicity theorems of Topkis (1978) and Milgrom and Shannon (1994) to analyse the responses of firms and industries to exogenous industry-wide shocks. Changes in the competitive environment are shown to be crucial for the comparative statics. Building on Hopenhayn (1992) and Melitz (2003), we put forward a model of monopolistic competition in which heterogeneous firms each make a decision on a number of activities. An activity refers to any variable at the discretion of the firm. Firms endogenously enter and exit the industry. The demand level of the industry inversely reflects the intensity of competition and changes to ensure a zero expected value of entry. While in fact endogenous in the aggregate, individual firms perceive the demand level as exogenous. Importantly, the activities faced by firms are complementary with: (i) each other; (ii) firm productivity; (iii) the demand level; and (iv) the exogenous industry-wide

¹See e.g. Milgrom and Roberts (1990a, 1995), Milgrom et al. (1991), Holmstrom and Milgrom (1994), Athey and Schmutzler (1995), Topkis (1995), and Milgrom and Roberts (1996).

parameters that we vary. Our focus lies on the implications of these complementarities which have previously been microfounded in a variety of ways.² An advantage of this approach is that the resulting model allows for quite general functional forms for the distribution of firm productivity and demand structures such as the additive, quadratic, and translog.

Since Melitz (2003), models with monopolistic competition and heterogeneous firms have become all the rage in international trade. The present paper provides a unifying framework for a vast number of these trade models. This framework is sufficiently general to encompass, at least symmetric-market versions of, well known contributions such as Melitz (2003), Antràs and Helpman (2004), Helpman et al. (2004), Melitz and Ottaviano (2008), Arkolakis (2010), Helpman and Itskhoki (2010), Helpman et al. (2010), Bernard et al. (2011), Bustos (2011), Davis and Harrigan (2011), Caliendo and Rossi-Hansberg (2012), Amiti and Davis (2012), Arkolakis et al. (2012), Antràs et al. (2014), Forslid et al. (2014), and Mayer et al. (2014). Despite large differences in focus (spanning vertical integration, FDI, markups, advertising, screening of workers, multiproduct firms, technology upgrading, efficiency wages, firm hierarchy, CO_2 emissions, and various types of international trade), these nested trade models share some important traits and a similar mathematical structure. By focusing on these common traits—including the assumptions on complementarities—and employing monotonicity theorems, we contribute by deriving strong, novel, and testable MCS predictions.

At the firm level, we investigate how the decisions of individual firms respond to exogenous increases in industry-wide parameters of the profit maximisation problem. Examples could be a decrease in the cost of undertaking (higher levels of) a given activity or the advent of a new activity that becomes available to all firms. These parameter changes do not only have a nonnegative direct effect (given the demand level) on the equilibrium decisions of individual firms as in standard models but also an indirect effect operating through induced changes in the demand level. Only when the indirect effect is nonnegative and thus aligned with the direct effect, we can be sure that the firm-level comparative statics are monotone. As a first result we provide sufficient conditions for MCS at the firm level by restricting the nature of the exogenous changes such that the demand level increases. The indirect

²Topkis (1995) and Mrazova and Neary (2013) consider conditions for complementarities to arise and provide examples. The models of Bustos (2011) and Amiti and Davis (2012) are very illustrative in relation to the assumed complementarities.

effect is, however, generally ambiguous and, as a consequence, so are the comparative statics at the firm level. This result illustrates how an endogenous competitive environment refines the comparative statics obtained in Milgrom and Roberts (1990a, 1995), Milgrom et al. (1991), Holmstrom and Milgrom (1994), Athey and Schmutzler (1995), and Topkis (1995). While MCS at the firm level are in fact also possible when the competitive environment is endogenous, the induced change in competition is decisive. We show that it is often straightforward to determine the direction of change in competition.

Our main finding is that firm-level complementarities may manifest themselves more clearly in the comparative statics for the industry composition than at the firm level. To see this possibility, which is particularly relevant for the nested models of international trade, let us first define MCS for the industry composition. By this, we mean that exogenous increases in industry-wide parameters of the profit maximisation problem lead to first-order stochastic dominance (FSD) shifts in the equilibrium distribution of any activity across firms. This implies that the share of active firms undertaking at least a given level of any activity increases and so does the average level of any activity. The main finding involves MCS for the industry composition when the comparative statics at the firm level are ambiguous. This possibility occurs when log-productivity is distributed across firms with nonincreasing hazard rate while the exogenous industry-wide shock enhances competition. This finding, for which the intuition is provided below, is important because: (i) trade models often analyse (trade) shocks that enhance competition; (ii) the common assumption of Pareto-distributed productivities implies that log-productivity is distributed with constant hazard rate.³ Testing the nested trade models at the industry level is therefore a promising empirical strategy. For instance, we show that trade liberalisation leads to FSD shifts in the firm-size distribution in a large subset of the nested trade models.⁴

A related main finding is that firm-level complementarities manifest themselves equally clearly in the comparative statics for the industry composition and at the firm level when log-productivity is distributed with increasing hazard rate. This holds true when productivity is log-normally or Frechet distributed. In this case, MCS appear at both the firm level and for the

³Much of the work after Melitz (2003) has applied the Pareto distribution for both reasons of tractability and also its empirical support provided by, among others, Axtell (2001), Luttmer (2007), and more indirectly by Eaton et al. (2011).

⁴This prediction holds even though some firms choose to become smaller after the trade liberalisation and its induced enhancement of competition.

industry composition when competition is dampened, while our results are ambiguous at both levels of analysis when competition is enhanced. From an aggregation standpoint, this finding seems a priori more intuitive than the opposite case with a nonincreasing hazard rate of log-productivity. To summarise our industry-level findings: the interaction between the hazard rate of log-productivity and the induced effect on competition is decisive for the comparative statics. Because of important selection effects via entry and exit of firms, MCS at the firm level are neither necessary nor sufficient for MCS for the industry composition. The paper further shows that a Pareto distribution of firm productivity assures MCS for the industry composition regardless of the induced effect on competition. This is due to a special knife-edge property.

Our results relate to the ongoing discussion about the actual distribution of firm productivity. Recently, Combes et al. (2012) and Head et al. (2014) have provided evidence in favour of the log-normal. Both groups of researchers show that a log-normal distribution of firm productivity provides a better description of their data than a Pareto distribution. Our theoretical results can be exploited empirically to further gauge the performance of, for instance, these two distributions. Let us provide an example. Using the model of Arkolakis et al. (2012), we show that an increase in market size makes all firms decrease their relative markup (an activity) over marginal cost. The distribution of markups across firms is however entirely unaffected as is the average markup if and only if productivity is Pareto distributed. The reason for this knife-edge outcome is a precisely offsetting selection effect through exit of low-productivity firms charging low markups. If instead productivity is log-normally distributed, the distribution of markups across firms shifts to the left in line with the pro-competitive effect at the firm level. Now, the selection effect is dominated by the pro-competitive effect at the firm level. Finally, if log-productivity is distributed with decreasing hazard rate, then we see MCS for the industry composition, wherefore the average markup increases due to a dominating selection effect.

Previously, Bernard et al. (2003) and Arkolakis et al. (2012) have also noted the discrepancy between firm- and industry-level effects of trade liberalisation on firm markups. These authors also use models with heterogeneous firms and endogenous selection. Yet, to the best of our knowledge, the present paper is the first to provide a general and thorough analysis of how firm-level complementarities can imply MCS for the industry composition despite ambiguities in firm-level responses. Further, we contribute by

illustrating how the industry-level implications depend on the distribution of firm productivity and departures from the common assumption of Pareto-distributed productivities. The paper perhaps closest to ours is the paper by Mrazova and Neary (2013). These two authors emphasise the role of supermodularity (complementarity) in shaping the sorting pattern of firms in a given equilibrium. Our approach differs by not only focusing on a given equilibrium but rather conducting comparative statics across equilibria. Costinot (2009) examines the role of log-supermodularity in generating comparative advantage. In the working-paper version, Costinot (2007) considers applications to specific heterogeneous firms setups. Here, comparative statics with respect to the productivity dispersion are conducted, whereas we consider comparative statics with respect to parameters directly affecting the maximisation problem of the firms. Finally, our paper relates to the studies of games with strategic complementarities or substitutes in industrial organisation where monotonicity theorems have also been applied; see for instance Topkis (1979), Jeremy et al. (1985), Milgrom and Roberts (1990b), Vives (1990), and Amir (2005). In contrast to these studies, which also allow for indirect effects through changes in the competitive environment, we rely on the simplified interaction among firms resulting from monopolistic competition. The key role of the demand level in our study also means that our study shares certain traits with the analysis of aggregative games in Acemoglu and Jensen (2013). Let us emphasise that, although the present paper has very close ties to the international-trade literature, nothing in the formulation of our framework limits the relevance or application of our results to trade-related issues.

The remainder of the paper is organised as follows. Section 2 develops our model and presents our central assumption of complementarities. Section 3 derives sufficient conditions for MCS at both the firm and the industry levels of analysis. Section 4 presents a concrete application of both our model and our results. Section 5 offers some concluding remarks. Appendix A briefly reviews the central mathematical results from Topkis (1978, 1995) and Milgrom and Shannon (1994) that we draw upon in our analysis as these may be unfamiliar to some readers.

2 Model

After paying a sunk entry cost of f_e units of the numéraire, atomistic firms enter an industry characterised by monopolistic competition. Upon entry, a firm realises its productivity level, $\theta \in [\theta_0, \infty)$ where $\theta_0 \geq 0$. Individual firms are fully characterised by their productivity level, θ , which is a realisation of the continuous random variable Θ with c.d.f. $F(\Theta)$.⁵ We let $F(\Theta)$ be strictly increasing on its entire support on which it is C^1 . Firms with strictly higher θ are assumed to be able to earn strictly higher profits.

After realising its productivity level, a firm has to choose whether to start producing or to exit the industry. If a firm chooses to produce, it has to make a decision, $x = (x_1, \dots, x_n)$, where x_i denotes the chosen level of activity i . An activity can refer to any variable at the discretion of the firm. The level of an activity can be either discrete or continuous. We let $x \in X$ where $X \subseteq \mathbb{R}^n$ is the set of all conceivable, but not necessarily available, decisions. The set X is assumed to be a lattice which, loosely speaking, means that undertaking a higher level of any activity may require, but importantly cannot prevent, undertaking a higher level of another activity. Restricting attention to lattices will allow complementarities between the n activities in x to take effect. The profitability of the decisions in X is influenced by a vector of exogenous industry-wide parameters, $\beta \in B$, with $B \subseteq \mathbb{R}^m$. Further, the actual choice set of all firms is restricted to a set of available decisions, $S \subseteq X$, with S being a sublattice of X . Our comparative statics will focus on changes in (β, S) which determines the attractiveness (all else equal) and availability of activities.

The Melitz (2003) model of international trade conforms to our setup and provides an example of activities, parameters, and choice sets. In this case, the decision, x , could comprise the export status and the amount of labour to employ, while β would contain the (negative) fixed and variable costs of exporting. A move from autarky to costly trade is an example of increasing the choice set of firms, S . Section 3.5 will later exploit this example to illustrate our results.

⁵In Section 3.8, we extend the model to allow for multidimensional firm heterogeneity. θ could in principle represent any firm characteristic that conforms to our assumptions.

2.1 The Demand Level

Firm profits depend on a single, common, and endogenous aggregate statistic which captures the (inverse) level of competition in the industry. We will refer to this variable, $A \in \mathbb{R}_+$, as the demand level and let firm profits be strictly increasing and continuous in A . In line with monopolistic competition among atomistic firms, individual firms perceive A as exogenous.⁶ To get a sense of what A could be, consider a model where the consumers' preferences are additively separable across varieties of a differentiated good. In this case, the inverse marginal utility of income enters the profit function through the demand function as a demand shifter and constitutes the demand level. Details will be provided in Section 3.7. A can potentially also comprise other endogenous variables such as factor prices as long as all endogenous variables outside the control of the firm can be combined into the single demand level A . That firm profits depend on just a single demand level is important to bear in mind. This implies that, while our model encompasses the model of Melitz (2003), it does not encompass the generalisation to asymmetric countries in Melitz and Redding (2014). This is because the profits of exporters depend on different demand levels (one for each export destination) when countries are asymmetric in Melitz and Redding (2014).⁷ However, country asymmetry does not necessarily pose problems. Two cases in point are the models by Antràs and Helpman (2004) and Antràs et al. (2014) which are encompassed by the model below.

2.2 Profits, Complementarities, and the Optimal Decision

Profits, π , of a firm with productivity θ depend on the decision, x , the demand level, A , and the industry parameters, β . We assume throughout that A and θ only enter the profit function through their product, $A\theta$. This assumption is very often satisfied in models of heterogeneous firms and discussed thoroughly in both Section 3.7, which shows that the assumption does not depend on CES preferences, and Appendix D. Formally,

$$\pi = \pi(x; A\theta, \beta), \tag{1}$$

⁶This setup also encompasses the case of perfect competition. To see this, let all firms share the same θ , let $f_e = 0$, and let A be the endogenous price level. For our industry-level analysis to be interesting, firm heterogeneity is however central.

⁷Results about MCS for the industry composition can also be obtained in models where firm profits depend on different demand levels. The analysis is available upon request.

where the semicolon separates choice variables from arguments that are perceived as exogenous by the firms. The following assumption summarises the three key complementarities in our model.

Assumption 1. *For all $(A\theta, \beta)$, the profit function, $\pi(x; A\theta, \beta)$, is supermodular in x on X and exhibits increasing differences in $(x, A\theta)$ and (x, β) on $X \times \mathbb{R}_+$ and $X \times B$, respectively.*

Supermodularity in x implies that the n activities are complementary. Milgrom and Roberts (1990a) argue that strategy and structure in modern manufacturing firms reflect widespread complementarities among many diverse activities in e.g. marketing, manufacturing, engineering, design, and organisation. In such a context, an increase in the choice set, S , can e.g. capture the advent of a new marketing technique that becomes available to all firms. The implied effect on the level of competition in our model is captured by an endogenous change in A . Topkis (1995) also considers the conditions for complementarities to arise and provides examples. Further, an extensive body of recent research within international trade relies heavily on complementary activities. Parts of this literature are surveyed in Section 9 in Melitz and Redding (2014). For illustrative examples, see for instance Bustos (2011), where technology upgrading and exporting are complementary activities, and Amiti and Davis (2012) where offshoring and exporting are complementary activities. The assumption of increasing differences implies that productivity, the demand level, and the elements of β are all complementary to the n activities.⁸

Proper ordering of activity levels and parameters is crucial for profits to satisfy Assumption 1.⁹ Even after proper ordering, Assumption 1 may not apply to all conceivable activities that firms face. However, if one can express the decision of firms in a form where the corresponding profit function satisfies our assumptions, then our results can be applied to the activities that constitute that decision.¹⁰ Consequently, we do not necessarily require

⁸Note that β only contains those parameters that comply with Assumption 1. As other parameters are kept constant throughout, we simply abstract from these.

⁹If a function is supermodular in (x_1, x_2) , then it is not supermodular in $(-x_1, x_2)$. If a function has increasing differences in (x_1, x_2, β) , then it does not have increasing differences in $(x_1, x_2, -\beta)$.

¹⁰For example, in the model of Helpman et al. (2004), exporting and FDI are not complementary. Once we recognise that exporting and FDI are not two different activities, but rather two different choices for a single activity concerning foreign market access, this

that all possible activities faced by firms are complementary. Lastly, note that profits in (1) do not have to represent a certain payoff to firms. In case of uncertainty after a firm has realised its productivity level and made its decision, (1) could be interpreted as expected profits; see e.g. Athey and Schmutzler (1995).

Faced with the profit function (1), a firm makes its optimal decision, x^* , under the constraint that $x \in S$, while taking θ , A , and β as given. Formally we have that

$$x^*(A\theta, \beta, S) = \arg \max_{x \in S} \pi(x; A\theta, \beta).$$

Lemma 1. *The optimal decision, $x^*(A\theta, \beta, S)$, is nondecreasing in $(A\theta, \beta, S)$.*

Lemma 1 follows readily from Theorem 1 in Appendix A and is simply the manifestation of the three key complementarities of Assumption 1. Importantly, these comparative statics are partial in nature since the endogeneity of A is ignored which will prove to be important. The profits obtained under the optimal decision are defined as

$$\pi^*(A\theta, \beta, S) \equiv \max_{x \in S} \pi(x; A\theta, \beta).$$

2.3 Entry

Firm profits upon entry are bounded below by zero because the firm exits the industry and forfeits the sunk cost of entry when π^* happens to be negative. Expected profits upon entry are thus given by

$$\Pi(A, \beta, S) \equiv \int \max\{0, \pi^*(A\theta, \beta, S)\} dF(\theta),$$

and are assumed to be finite. It is well known that this may require some restrictions on the distribution of productivities; see Melitz (2003). We assume unrestricted entry and an unbounded pool of potential entrants. In equilibrium, the expected profits upon entry must therefore be equal to the cost of entry,

$$\Pi(A, \beta, S) = f_e. \tag{2}$$

model can be analysed within the present framework. Further, Milgrom et al. (1991) discuss how to obtain a supermodular profit function in a specific case with a core group of complementary activities and a group of additional activities.

Given the existence of an equilibrium value of A that satisfies (2), uniqueness is ensured by profits being strictly increasing in A . Existence is ensured by continuity and by assuming that firms cannot earn strictly positive profits as demand vanishes ($\lim_{A\theta \rightarrow 0} \pi^* \leq 0$) and that the demand level can become sufficiently high for firms to be able to recoup the entry cost ($\lim_{A\theta \rightarrow \infty} \pi^* > f_e$). Then (2) pins down the endogenous demand level as a function of β , S , and f_e .

2.4 Industry Composition

We denote the c.d.f. of the equilibrium distribution of activity i across active firms by $H_i(x_i; \beta, S)$, $i = 1, \dots, n$. These distributions are the focus of our industry-level analysis in Sections 3.2 and 3.3. To characterise these distributions, consider the cross-section of firms in a given equilibrium where (A, β, S) is given. Since firms with strictly higher θ are able to earn strictly higher profits, the self-selection or sorting of firms into being active or exiting obeys the rule that all firms with productivities above a certain threshold are active and all firms with productivities below exit. Denoting this threshold by θ_a , we have that

$$\theta_a(A, \beta, S) \equiv \inf\{\theta : \pi^*(A\theta, \beta, S) > 0\}.$$

We focus on the case with endogenous exit by assuming that the lowest productivity firms are not able to produce profitably. That is, $\theta_a(A, \beta, S) > \theta_0$. The underlying reason could e.g. be the presence of some fixed costs of production or the presence of a choke price. The next step is to characterise the sorting of active firms into the activities based on productivity. By Lemma 1, the complementarities of our model imply that, in a given equilibrium, higher productivity firms choose weakly higher levels of all activities which echoes one of the main points made by Mrazova and Neary (2013). Let θ_i be the lowest level of productivity at which a firm undertakes at least level x_i of activity i . Bounding this threshold from below by θ_a , it is given by

$$\theta_i(x_i; A, \beta, S) \equiv \max\{\theta_a, \inf\{\theta : x_i^*(A\theta, \beta, S) \geq x_i\}\}. \quad (3)$$

On the basis of the above sorting pattern, we now characterise the equilibrium distributions of activities and the industry composition. In the following, we focus on a particular level, x_i , of activity i , which could be any level of any of the n activities. Applying the law of large numbers, let $s_a \equiv 1 - F(\theta_a)$

be the share of firms that are active and let $s_i \equiv 1 - F(\theta_i)$ denote the share of firms undertaking at least level x_i of activity i . Note that $s_a \geq s_i$. Using these shares, the c.d.f. of the equilibrium distribution of each activity i can be expressed as

$$H_i(x_i; \beta, S) = 1 - \frac{s_i(x_i; A(\beta, S), \beta, S)}{s_a(A(\beta, S), \beta, S)}. \quad (4)$$

The industry composition refers jointly to these n distributions.

3 Comparative Statics

We now investigate the equilibrium responses of individual firms and of the industry composition to increases in the industry-wide parameters (β, S) . Consistent with Lemma 1, both increases in β and increases in S provide firms with an incentive to increase their levels of all activities, all else equal. Increases in β do so by increasing the attractiveness of undertaking higher levels of the activities, while increases in S do so by shifting upwards the choice set of available decisions.¹¹ Importantly, these incentives can be brought about in two distinct ways. By increasing β , the attractiveness of higher levels of activities increases both if profits associated with higher levels increase and if profits associated with lower levels decrease. Analogously, S is shifted upwards both if higher levels of activities become available and if lower levels become unavailable. While these two types of increases in β and S are not mutually exclusive, their distinct effects on firm profits will prove to be crucial for the comparative statics.

Definition 1. *An increase in (β, S) is competition enhancing if expected profits upon entry, Π , increase given the demand level. If expected profits upon entry decrease given the demand level, the increase in (β, S) is competition dampening.*

The two terms competition-enhancing increases and competition-dampening increases in (β, S) simply refer to the equilibrium effect on the demand level, A . To see this, note that A responds to offset the direct effect of the change

¹¹ S can e.g. increase by allowing higher levels of existing activities. This obviously includes the case of allowing levels higher than zero of a given activity, i.e., introducing new complementary activities by changing S from S' to S'' such that $S' \subset S''$ and $S' \leq_s S''$.

in (β, S) on Π . On the one hand, if Π tends to increase given A , as a consequence of increasing (β, S) , A falls in order to satisfy (2). Firms perceive this decrease in A as enhanced competition. On the other hand, if Π tends to fall given A , the result is an increase in A and hence a dampening of competition from the firms' point of view. As profits depend on a single demand level, it is often straightforward to determine whether a change in (β, S) is competition enhancing or competition dampening. For example, reducing the costs of existing activities (increasing β) or introducing a new activity into S (increasing S) corresponds to a competition-enhancing increase in (β, S) . This is because the profits of all firms are at least weakly increasing in this exogenous variation given A . Such an argument explains why trade liberalisations in Melitz (2003) through decreases in either the fixed or variable costs of international trade (or a move away from autarky through an introduction of exporting) are competition enhancing. Consequently, the exact nature of the increase in (β, S) determines the direction of change in A .

3.1 Firm Level

Starting at the firm level, let us define the equilibrium decision of a firm conditional on being active as

$$\tilde{x}^*(\theta, \beta, S) \equiv x^*(A(\beta, S)\theta, \beta, S). \quad (5)$$

From the RHS of (5), it is clear that changes in (β, S) have a direct effect on firm decisions for a given demand level but, since the whole industry is affected, such changes also have an indirect effect through changes in the demand level. This dichotomy allows us to decompose the total effect on \tilde{x}^* from changing (β', S') to (β'', S'') where either β or S could remain unchanged. Note that

$$\Delta\tilde{x}^* = \underbrace{x^*(A'\theta, \beta'', S'') - x^*(A'\theta, \beta', S')}_{\text{Direct effect}} + \underbrace{x^*(A''\theta, \beta'', S'') - x^*(A'\theta, \beta'', S'')}_{\text{Indirect effect}},$$

where $A' = A(\beta', S')$ and $A'' = A(\beta'', S'')$. It follows from Lemma 1 that an increase in (β, S) always has a nonnegative direct effect on the equilibrium decision, \tilde{x}^* . The increase in (β, S) provides firms an incentive to increase their levels of at least one activity. The inherent complementarities among activities ensure that this is manifested in an increase in \tilde{x}^* , all else equal. Whereas the direct effect of an increase in (β, S) on \tilde{x}^* is unambiguously nonnegative irrespective of how (β, S) increases, recall that the sign of

the indirect effect critically depends on whether competition is enhanced or dampened. By Lemma 1, the sign of the indirect effect is equivalent to the sign of the change in A . Thus, the indirect effect is aligned with the direct effect when competition is dampened but opposed to the direct effect when competition is enhanced. The following proposition summarises.¹²

Proposition 1. *The total effect on the equilibrium decision, \tilde{x}^* , is nonnegative for all firms if the increase in (β, S) is competition dampening. If the increase in (β, S) is competition enhancing, the total effect on the equilibrium decision is ambiguous.*

The feedback from the endogenous demand level implies that the decisions of individual firms do not generally exhibit MCS in (β, S) in spite of the complementarities imposed by Assumption 1. The reason is that our assumption of complementarities is partial in nature and by no means ensures MCS at the firm level once the endogeneity of A is recognised. Importantly, the details of the increase in (β, S) thus matter which contrasts the case where the competitive environment is exogenous. New or better opportunities that become available to all firms can make a given firm scale down existing activities even when these existing activities are complementary to the activities affected. Thus, for some firms, such opportunities may turn out to be a threat detrimental to many dimensions of the firms' operations. This result still holds when F is degenerate such that all firms share the same θ and make the same decision, \tilde{x}^* , although one qualification must be given. In this case, an increase in S will unambiguously increase \tilde{x}^* .¹³ The total effect of an increase in β remains ambiguous when firms are homogeneous. A good example, which fits into our framework and this discussion, is the Krugman (1979) model. As shown by Zhelobodko et al. (2012), the total output of firms (an activity) may either increase or decrease following an increase in

¹²Note that the results in Proposition 1 are conditional on a firm being active. Hence, Proposition 1 does not rely on the assumption that $\theta_a > \theta_0$. Nor does it rely on the assumption of profits only depending on A and θ through $A\theta$ or the complementarity between x and θ . However, these assumptions play important roles for our industry-level analysis.

¹³To see this, first note that if the initial \tilde{x}^* becomes unavailable, then only higher decisions are available. Second, if the initial \tilde{x}^* remains available, we know, by the definition of the strong set order (\leq_s) in Section A, that lower decisions available ex post were also available ex ante. It follows that such decisions do not constitute an equilibrium after S has increased.

market size (β) depending on how the so-called relative love of variety varies with consumption.

One can obviously also use the approach above to derive some unambiguous results. The comparative statics are for instance monotone when the attractiveness of higher levels of an activity increases (β increases) because lower levels of this activity are taxed (and the tax revenue is not redistributed among firms in the industry). To see this, note that this increase in β induces an increase in A since the tax decreases the profits of all firms given A . Furthermore, it is often the case that some firms are not directly affected by an increase in (β, S) implying that the effect on the decision of these firms is straightforward to determine. One case in point is trade liberalisation in the many models following Melitz (2003) where the least productive active firms do not select into trading activities because of various fixed costs. In such a case, the least productive active firms lower their decision, \tilde{x}^* , as competition enhances. We provide an example of this in Section 3.5.

3.2 Industry Level: a First Glance

We now move on to investigating how the industry composition responds to changes in (β, S) . Apart from the effects highlighted in our firm-level analysis above, selection effects through changes in θ_a are central for our industry-level analysis. These arise since some marginal firms may leave the industry or become active producers as a result of the change in (β, S) .

Our notion of monotone comparative statics (MCS) for the industry composition is formalised as follows.

Definition 2. *The industry composition exhibits MCS when increases in (β, S) induce first-order stochastic dominance (FSD) shifts in the equilibrium distributions of all activities. That is, $H_i(x_i; \beta, S)$ is nonincreasing for all levels, x_i , of all activities, $i = 1, \dots, n$.*

MCS for the industry composition thus mean that the equilibrium distributions of the n activities unambiguously shift towards higher values (such that the share of active firms that undertake at least any positive level of any activity increases). Consequently, the average level of any activity increases. Recall that an FSD shift in the equilibrium distribution of an activity implies FSD shifts in all firm-level variables that are monotonically increasing in this activity.¹⁴

¹⁴We thank an anonymous referee for pointing this out.

In order to build intuition for our industry-level results below, consider an increase in β and assume now for simplicity that H_i is differentiable in a scalar β . Then we can express the total effect of increasing β on H_i as¹⁵

$$\frac{1}{1 - H_i} \frac{dH_i}{d\beta} = \underbrace{-\frac{1}{s_i} \frac{\partial s_i}{\partial \beta} + \frac{1}{s_a} \frac{\partial s_a}{\partial \beta}}_{\text{Direct effects}} - \underbrace{\left(\frac{1}{s_i} \frac{\partial s_i}{\partial A} - \frac{1}{s_a} \frac{\partial s_a}{\partial A} \right) \frac{dA}{d\beta}}_{\text{Indirect effects}}. \quad (6)$$

Given the assumption of endogenous exit ($\theta_a > \theta_0$), the equilibrium distributions, H_i , are both affected by the levels of the activities undertaken by firms conditional on being active (level effect) and by the endogenous selection of which firms are active (selection effect). In (6), the level effect is represented by the effects on the share of firms undertaking at least level x_i of activity i , s_i , while the selection effect is represented by the effects on the share of firms that are active, s_a . Each of these effects has a direct component, which is the effect for a given A , and an indirect component, which is the effect through a change in A . The total level effect corresponds to the firm-level responses analysed in Section 3.1.

The decomposition for the general case, which is covered by Propositions 2-4, is presented in Appendix C. In the general case, we do not assume differentiability of H_i with respect to a scalar β , and the decomposition is valid for changes in S as well. Anyway, (6) is more elegant and sufficient to illustrate how the total selection effect introduces a possible discrepancy between the firm-level comparative statics (the total level effect) and the comparative statics for the industry composition. Absent a selection effect, the comparative statics for the industry composition follow directly from the firm-level comparative statics. In this case, the industry composition exhibits MCS (in general) if and only if the total effect on \tilde{x}^* is nonnegative for all firms. However, the trade literature almost exclusively focuses on the arguably more interesting and appealing case where some firms endogenously shut down or enter after a change in θ_a . As we shall see, MCS at the firm level are neither necessary nor sufficient for MCS for the industry composition in this situation.

3.3 Sufficient Conditions for MCS for the Industry Composition

In order to derive sufficient conditions for MCS for the industry composition, start by considering the direct level and selection effects on H_i . In the

¹⁵We thank J. Peter Neary for suggesting the representation (6) to us.

simplified case of differentiability, these correspond to the first and second term on the RHS of (6), respectively. For the more general case, consult Appendix C. Since an increase in (β, S) tends to increase the levels of the activities chosen by individual firms (for a given A), it tends to increase s_i .¹⁶ The direct level effect on H_i is therefore nonpositive which works in favour of MCS for the industry composition. For the direct selection effect to be nonpositive as well, we simply need the direct effect on s_a to be nonpositive. Intuitively, the marginal active firms have low productivities and therefore undertake relatively low levels of the activities conditional on being active. An increase in the share of active firms therefore works against MCS for the industry composition. Whether the direct effect of an increase in (β, S) on s_a is indeed nonpositive needs to be checked when using our approach. For now it suffices to point out that this pivotal condition may constrain the nature of the increase in (β, S) . The condition will be thoroughly discussed after we have presented Proposition 2 which together with Proposition 3 depend on this pivotal condition.

Next, consider the two indirect effects on H_i , i.e., the indirect level and selection effects. These effects correspond to the entire last term on the RHS of (6). An increase in A tends to make all firms weakly increase their levels of activity i . This makes the indirect level effect on H_i nonpositive.¹⁷ At the same time, an increase in A will allow some previously inactive low-productivity firms to produce profitably. The indirect selection effect is therefore non-negative. These two indirect effects are reversed when A decreases but are obviously still opposing. Therefore, in order to ensure that the sum of these two indirect effects on H_i is nonpositive regardless of the possibly unknown direction of change in A , we must require that the indirect level and selection effects are exactly offsetting. This is the case if a change in A induces the same percentage of changes in s_a and s_i . It is obvious that the total indirect effect in (6) is zero in this case. The following proposition, which is proven in Appendix C, makes it clear that this requires a Pareto distribution, $F(\Theta)$, or firm productivities.

Proposition 2. *Increases in (β, S) induce MCS for the industry composition*

¹⁶To see this formally, note that by Lemma 1, x^* is nondecreasing in $(A\theta, \beta, S)$. Thus it follows from (3) that θ_i is nonincreasing in (β, S) given A . Therefore, $s_i = 1 - F(\theta_i)$ is nondecreasing in (β, S) given A .

¹⁷It follows from Lemma 1 and (3) that θ_i is nonincreasing in A given (β, S) . Therefore, s_i is nondecreasing in A given (β, S) .

if the direct effect on s_a is nonpositive and if log-productivity is distributed with constant hazard rate. The latter condition is equivalent with productivity being Pareto distributed.

This key result depends on two important conditions which will now be discussed. We start by discussing the condition that $F(\Theta)$ is Pareto. A constant hazard rate of log-productivity implies by definition that the density at any level of log-productivity is constant relative to the probability mass above it. This means that the percentage changes (induced by a change in A) in the share of active firms and the share of firms undertaking at least a given level of activity i are equal if the changes in the log-thresholds $\log \theta_a$ and $\log \theta_i$ are equal. But this is implied by profits depending on A and θ only through $A\theta$. As a consequence, the indirect level and selection effects exactly cancel out. Under nonpositive direct level and selection effects, this gives us MCS for the industry composition. While the Pareto distribution is often used in models of international trade due to other attractive features, Proposition 2 points out a novel knife-edge property with strong implications for industry-level comparative statics.

Proposition 2 is moreover conditional on a nonpositive direct effect on s_a . The intuition for the need of this pivotal condition, which assures a nonpositive direct selection effect and MCS for the industry composition given that $F(\Theta)$ is Pareto, was provided earlier. The condition implies that the direct effect of increases in (β, S) on the profits of the least productive active firms must be nonpositive such that θ_a is nondecreasing in (β, S) given A . It is important to note that, under this condition, increases in (β, S) can still be both competition enhancing and competition dampening. On the one hand, the condition is clearly satisfied when the direct effect on the profits of all firms is nonpositive. On the other hand, the condition is satisfied when the direct effect of an increase in (β, S) on the profits of the least productive active firms is zero while the direct effect on the profits of all other firms is nonnegative and positive for some. Importantly, this latter scenario is very often seen in models of international trade, such as those listed in the introduction, since the least productive active firms are very often not directly affected by the comparative statics considered. One case in point is competition enhancing trade liberalisations in models where the least productive firms do not select into trading activities because of various fixed costs these firms cannot overcome. We will provide a few examples of this pivotal condition, which is easily checked for by simple inspection of the

profit function, at later stages in this paper.

As argued in the beginning of Section 3, it is often straightforward to determine whether a change in (β, S) is competition enhancing or competition dampening. Hence, it can now be argued that both the exact nature of the increase in (β, S) and often also the selection into various activities (like e.g. exporting) jointly determine the direction of change in A and whether the direct effect on s_a is nonpositive as needed. Furthermore, once we consider increases in (β, S) that we know are either competition enhancing or competition dampening, we no longer need to be on the knife edge where the indirect level and selection effects on H_i exactly balance. This leads us to the following proposition which is also proven in Appendix C.¹⁸

Proposition 3. *Competition-enhancing increases in (β, S) induce MCS for the industry composition if the direct effect on s_a is nonpositive and if the distribution of log-productivity has a nonincreasing hazard rate. Competition-dampening increases in (β, S) induce MCS for the industry composition if the direct effect on s_a is nonpositive and if the distribution of log-productivity has a nondecreasing hazard rate.*

To understand the intuition behind Proposition 3, remember that A induces the indirect level and selection effects through its effects on θ_i and θ_a , respectively. Relative to the case with constant hazard rate of log-productivity (i.e., the knife-edge case where the two indirect effects balance), a nonincreasing hazard rate puts more relative probability density at $\log \theta_a$ relative to $\log \theta_i$ since $\theta_a \leq \theta_i$. Consequently, the indirect selection effect dominates the indirect level effect. This works in favour of MCS for the industry composition when A falls. Conversely, a nondecreasing hazard rate of log-productivity means that the indirect level effect dominates, which works in favour of MCS for the industry composition when A increases. Relating to the case of differentiability in (6), the hazard-rate conditions of Proposition 3 ensure that $(\frac{1}{s_i} \frac{\partial s_i}{\partial A} - \frac{1}{s_a} \frac{\partial s_a}{\partial A})$ and $\frac{dA}{d\beta}$ share the same sign, wherefore the indirect effects are nonpositive in total. Taking stock of this section's insights, we see that Proposition 2 is useful since it reveals sufficient conditions under which MCS for the industry composition can be obtained in the vast number

¹⁸Distributions with monotone hazard rates include, among others, the Gumbel, the exponential, the Weibull, the gamma, and the normal. The normal and Gumbel distributions exhibit monotone increasing hazard rates. Recall that log-productivity is distributed Gumbel when productivity is Frechet distributed as in Eaton and Kortum (2002) and Bernard et al. (2003).

of trade models that fit into our framework. Proposition 3 shows that the same results can often be obtained under other distributional assumptions than just the Pareto. Note that one does not need to solve a full general equilibrium model to get to the results above.

3.4 Firm-Level versus Industry-Level Results

Having described the comparative statics both at the firm level and for the industry composition, we summarise our findings in Table 1. The illustrated effect on H_i is conditional on a nonpositive direct effect on s_a . The hazard rate of the distribution of log-productivity is denoted by $\lambda_{\log \theta}$.

Comparative Statics for an Increase in (β, S)				
Effect on competition	Effect on \tilde{x}^*	Effect on H_i		
		$\frac{d\lambda_{\log \theta}}{d \log \theta} \geq 0$	$\frac{d\lambda_{\log \theta}}{d \log \theta} = 0$	$\frac{d\lambda_{\log \theta}}{d \log \theta} \leq 0$
Dampened	Increasing	FSD	FSD	?
Enhanced	?	?	FSD	FSD

Table 1: Summary of firm- and industry-level comparative statics.

These results point out a clear discrepancy between the conditions that ensure MCS at the firm level and those that ensure MCS for the industry composition. At the firm level, MCS hinge upon the particular shock considered, i.e., whether competition is enhanced or dampened. Under Pareto-distributed productivities, the industry composition exhibits MCS regardless of the effects on competition (provided the direct effect on s_a is nonpositive). Thus, firm-level complementarities may manifest themselves much more clearly in the comparative statics for the industry composition than at the firm level. This possibility occurs when the increase in (β, S) enhances competition as is often the case. Tests of firm-level complementarities conducted at the industry level may therefore be a promising empirical strategy for testing the nested trade models with Pareto-distributed productivities. More generally, we see that the interaction between the shock type and the properties of $\lambda_{\log \theta}$ is decisive at the industry level when the direct effect on s_a is nonpositive. If log-productivity is distributed with decreasing hazard rate, we are ensured MCS for the industry composition when firm-level responses are ambiguous (competition is enhanced). Under the same circumstances,

we see ambiguous shifts in the industry composition when the firm-level responses are unambiguous (competition is dampened). Nontrivial selection effects hence imply that MCS at the firm level are neither necessary nor sufficient for MCS for the industry composition. It is entirely possible that firm-level responses to an exogenous shock to the industry are ambiguous and depend on the characteristics of the firms and activities in question while the industry composition unambiguously shifts towards higher levels of all activities. Notice from columns 2 and 3 in Table 1 that MCS for the industry composition and MCS at the firm level arise in the same case (competition is dampened) when productivity is distributed, e.g., log-normally or Frechet, since this implies an increasing hazard rate of log-productivity.¹⁹ From an aggregation standpoint, this seems more appealing than the case with a non-increasing $\lambda_{\log \theta}$, but determining the relevant case is of course an empirical question.

3.5 Trade Liberalisation and the Firm-Size Distribution

To illustrate our firm- and industry-level results, consider as an example from international trade a two-country Melitz (2003) model. We let the activities of the firms be export status, given by the indicator $\mathbb{1}_{ex}$ for exporting, and total labour input for variable production, l . If we require that the resulting output is optimally distributed across markets in case of exporting, then we obtain the profit function

$$\pi(l, \mathbb{1}_{ex}; A\theta, \beta) = (1 + \mathbb{1}_{ex}\tau^{1-\sigma})^{\frac{1}{\sigma}} l^{\frac{\sigma-1}{\sigma}} (A\theta)^{\frac{\sigma-1}{\sigma}} - l - f - \mathbb{1}_{ex}f_{ex},$$

where $\sigma > 1$ is the elasticity of substitution, $\tau > 1$ is an iceberg trade cost, f is a fixed cost of production, and f_{ex} is a fixed cost of exporting. Assumption 1 is satisfied if we let $x = (l, \mathbb{1}_{ex})$, $X = \mathbb{R}_+ \times \{0, 1\}$, and $\beta = (-\tau, -f_{ex})$. Since the model conforms to our setup, we can apply our propositions to analyse the effects of a reduction in τ .²⁰ A reduction in τ is clearly competition

¹⁹This again requires the direct effect on s_a to be nonpositive. Note that, when productivity is distributed log-normally or Frechet, we also see ambiguous responses at both levels of analysis when competition is enhanced.

²⁰We could also consider an opening to trade and incremental liberalisations of trade through a reduction in f_{ex} . Opening to trade implies introducing a new activity (exporting) by moving from $S' = \mathbb{R}_+ \times \{0\}$ to $S'' = \mathbb{R}_+ \times \{0, 1\}$ which constitutes an increase in S . However, these two types of trade liberalisation have effects similar to the liberalisation through τ . Note that both types of trade liberalisation are competition enhancing.

enhancing since it weakly increases the profits of all firms given the demand level. Starting with the firm level, this implies that the direct and indirect effects on x are opposing. All exporters nevertheless increase their use of labour while (ex-post) nonexporters reduce their use of labour. While the positive direct effect dominates the negative indirect effect on the use of labour for exporters, nonexporters are only affected by the negative indirect effect.²¹

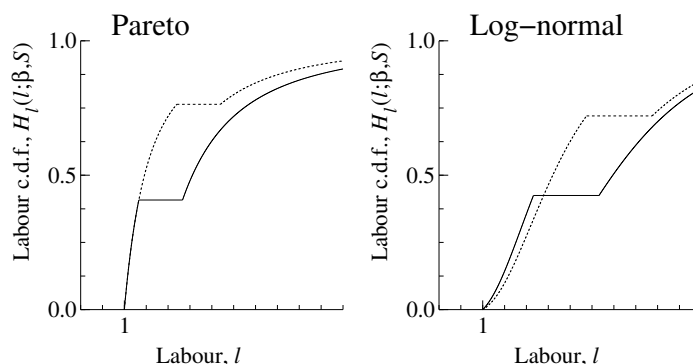


Figure 1: Equilibrium distribution of labour input before (dotted) and after (solid) a trade liberalisation depending on the productivity distribution.

Moving on to the industry level, note that the trade liberalisation has no direct effects on the profits of the least productive firms which do not export by assumption. This is assured through assumptions on the size of exogenous parameters that govern the cost of international trade. That is, the direct effect on s_a is nonpositive, and in fact zero. Thus, we know that if productivity is Pareto distributed, then the trade liberalisation induces a FSD shift in the firm-size distribution.²² However, this FSD shift in the firm-

²¹This shows that the direction of firm-level responses may vary across firms for a given activity. To see that responses can vary across activities within a given firm, one can split total labour for variable production into that used for production to the domestic market and the export market. Upon a decrease in τ , exporters decrease and increase their use of labour for production to the domestic and the export market, respectively.

²²This holds whether we use either total labour input for variable production or firm revenue as the measure of firm size when the trade liberalisation appears through a decrease in f_{ex} . The reason being that, in Melitz (2003), firm revenue is monotonically increasing in labour demand for variable production. The resulting increase in average revenue in turn implies that the number of active firms decreases.

size distribution does certainly not happen if productivity is log-normally or Frechet distributed since this means that log-productivity is distributed with strictly increasing hazard rate.²³ Figure 1 illustrates these points and also shows that the share of exporting firms increases (an FSD shift in the distribution of the exporting activity) under both distributional assumptions in this particular example. This can be seen from the downwards shift in the horizontal segments in Figure 1, located at the share of active firms which do not export.

So far we have showed that, in the Melitz (2003) model, trade liberalisations induce FSD shifts in the firm-size distribution as long as the distribution of log-productivities has nonincreasing hazard rate. Many subsequent extensions of the Melitz (2003) model expand the choice set of firms in a way such that the resulting model still features complementarities between labour input and the rest of the activities. Since the Pareto assumption is adopted in most of these studies, our results imply that trade liberalisations inducing FSD shifts in the firm-size distribution is an overarching prediction, which holds in models such as Helpman et al. (2004), Antràs and Helpman (2004), Arkolakis (2010), and Bustos (2011).²⁴ This prediction can thus be used to evaluate a whole strand of the recent heterogeneous-firms trade literature. If the prediction of FSD shifts in the firm-size distribution following trade liberalisations (or decreases in the costs of complementary activities) is not supported by data, one remedy could be to abandon the Pareto assumption in favour of a log-normal or Frechet assumption. As argued above, these distributions do not give rise to FSD shifts in the firm-size distribution provided that not all firms undertake the particular activities the attractiveness of which increase.

The distinct implications of productivity being either Pareto or log-normally distributed are interesting in the light of the on-going discussion about the distribution of firm productivities. Much of the work following Melitz (2003) has applied the Pareto distribution for its tractability and empirical support.²⁵ More recently however, Combes et al. (2012) and Head et al. (2014)

²³See the analysis in Appendix C where the direct level and selection effects are zero for a sufficiently low level of labour input while the total indirect effect is strictly positive.

²⁴Further, the prediction holds for Helpman and Itskhoki (2010), Helpman et al. (2010), Davis and Harrigan (2011), and Egger and Koch (2013) if one considers the distribution of expenditure on labour input instead of that of labour input itself.

²⁵See e.g. Axtell (2001), Luttmer (2007), and the more indirect evidence in Eaton et al. (2011).

have provided evidence in favour of the log-normal distribution. They argue that the log-normal distribution provides a much better description of their French and Chinese data relative to the Pareto distribution.^{26,27} These contrasting views concerning the right distributional assumption are noteworthy given the very different implications for comparative statics at the industry level pointed out above. We believe that our predictions for the firm-size distribution can be exploited to gauge the empirical performance of these two distributions which, in the present context, and in contrast to many other studies, are equally tractable because of our different mathematical approach.

3.6 Comparative Statics without Direct Effects

Some parameter changes have no direct effects on either firms' decisions or their choice to exit. While this could be true for some elements of β , it is obviously the case for the sunk entry cost, f_e , which we will use to illustrate the implications. First off, the firm-level comparative statics are determined solely by the indirect effect on \tilde{x}^* . By (2), it is clear that an increase in f_e must be competition dampening (A increases strictly) and thus increases the equilibrium decisions of all firms. At the industry level, things may be different. Since an increase in f_e does not have direct level or selection effects, only the indirect effects matter at the industry level. We can therefore conclude that the equilibrium distributions of the activities are completely unaffected by the increase in f_e if productivity is Pareto distributed. Further, if $\lambda_{\log \theta}$ is nondecreasing such that the indirect level effect dominates the indirect selection effect, then the equilibrium distributions of activities shift toward higher values when f_e increases. However, if $\lambda_{\log \theta}$ is nonincreasing such that the indirect selection effect dominates, then all of these distributions shift towards lower values regardless of the unambiguous increase in the equilibrium decisions of all firms. Based on (11) in Appendix C, it can also be argued that these distributional shifts to either the right or left are nontrivial when $\lambda_{\log \theta}$ is, respectively, strictly increasing or strictly decreasing. As these findings

²⁶Head et al. (2014) also argue that one has to be careful in defending the Pareto due to its performance in its right tail. In their paper, the left part of the tail is decisive for welfare changes due to selection effects which are also crucial for comparative statics in our model. These authors also mention the work of Eeckhout (2004) which shows that the log-normal and the Pareto distributions differ the most to the left.

²⁷See also Ijiri and Simon (1974).

are obviously also valid for the example in the previous section, it follows that an increase in the entry barrier, f_e , leads to a strict decrease in the number of active firms when $F(\Theta)$ is Frechet or log-normal. This intuitive finding does not necessarily hold when log-productivity is distributed with nonincreasing hazard rate.

3.7 Preference Structures with $A\theta$

Next, let us discuss preference structures where the $A\theta$ assumption may hold. Although CES preferences may be most commonly used—at least in models of international trade—we note that profits depending on θ and A only through $A\theta$ can arise under various other preference structures than CES.

First off, with a slight reinterpretation of θ as perceived quality, we show that multiplicative separability ($A\theta$) can arise under relatively general additively separable preferences. Let consumers have preferences $U = \int \theta_j u(c_j) dj$ with j indexing varieties of a differentiated good, while c_j and θ_j are the quantity consumed and the quality of variety j , respectively.²⁸ With identical consumers maximising utility subject to the budget constraint, $\int p_j c_j dj = I$, where p_j is the price of variety j and I is income, the inverse demand function reads $p_j = A\theta_j u'(c_j)$, with A being the inverse marginal utility of income. The revenue of a firm reads $r(q_j) = A\theta_j u'(q_j/L)q_j$, where q_j is total output and L is market size (the number of consumers). With this revenue function, the profit function is bound to depend on θ and A only through their product, $A\theta$.

Second, Arkolakis et al. (2012) consider markups under a general formulation of demand that encompasses three important groups of utility functions: additively separable (but non-CES), quadratic (as in Ottaviano et al., 2002; Melitz and Ottaviano, 2008), and translog. In a symmetric country version of their model, the profits of a firm in country i selling in market j can be written as

$$\pi_{ij}(x; A\theta, \beta) = L(\theta/\tau_{ij})^{\epsilon-1} (x-1)x^{-\epsilon} e^{d(\log x + \log \tau_{ij} - \log(A\theta))}, \quad (7)$$

where $\epsilon \leq 1$ is a parameter of the demand system, L is again the market size, τ_{ij} is an iceberg trade cost, and firms decide on their relative markup denoted

²⁸In this case, θ_j represents perceived quality in the sense that it scales up marginal utility, $u'(c_j)$, or marginal willingness to pay, for a given level of consumption.

by $x = p\theta/\tau_{ij}$, where p is the price. Further, d is a decreasing function, and our demand level, A , is the choke price in this case, i.e., setting a markup such that $p \geq A$ results in zero demand and hence zero profit.²⁹ Since the $\theta^{\epsilon-1}$ in front of the profit function neither affects the decision on activities nor the choice to exit, we can ignore this factor for all purposes except the free-entry condition. But this means that, for all relevant purposes, the profit function depends on θ and A only through $A\theta$. Let $X = [1, \infty)$.³⁰ Note that (7) exhibits the single crossing property in (x, L) and $(x, A\theta)$. Thus, we can apply our results with $\beta = L$.³¹ L does not affect the optimal decision, x^* , directly wherefore an increase in market size only has a negative indirect effect. To see this, note that for a given choke price (demand level) all firms can earn weakly higher profits the higher L is. Hence, a larger market ends up with a lower choke price, which makes all firms reduce their markups. This is the pro-competitive effect discussed by Arkolakis et al. (2012) and shown empirically by Bellone et al. (2014).

Next, consider the industry composition. The direct effect on s_a is non-positive under an increase in L since this increase does not directly affect the optimal profits of the least productive active firms which let $x = 1$ and earn zero profits. Thus, we can conclude that when productivities are Pareto distributed, an increase in L implies that the equilibrium distribution of markups and the average markup are completely unaffected by the increase in market size despite the pro-competitive effects at the firm level.³² If instead productivities are log-normally distributed, such that log-productivity is distributed with strictly increasing hazard rate, the equilibrium distribution of markups exhibits a nontrivial shift towards lower values in line with the pro-competitive effects at the firm level.³³ Finally, if the distribution of

²⁹This follows from the assumed properties of the function d ; see Arkolakis et al. (2012) for details.

³⁰No firm would want to choose a price lower than its marginal cost wherefore $x \geq 1$.

³¹As (7) is trivially supermodular in x , increasing differences in $(x, A\theta)$ and (x, β) is a sufficient condition for the partial comparative statics in Lemma 1. The single crossing property is both necessary and sufficient; see Appendix B.

³²This knife-edge balance between the indirect level and selection effects appears in Bernard et al. (2003) and Arkolakis et al. (2012). One of our contributions is to analyse the outcome once we leave this knife edge and relate industry-level comparative statics to the distributional assumptions, i.e., to derive distributional comparative statics (Jensen, 2014).

³³This relates to a remark made by Arkolakis et al. (2012). However, as our criterion is first-order stochastic dominance instead of monotone likelihood ratio dominance, our

log-productivity has a strictly decreasing hazard rate, then the equilibrium distribution of markups exhibits a nontrivial shift towards higher values. In this case, sufficiently many low-productivity firms are driven out of business by the drop in A to ensure higher markups on average when market size increases. Equivalent results, which accentuate the importance of formal aggregation, can be derived within a closed-economy Melitz and Ottaviano (2008) model.

Notice that these theoretical results concerning the average markup may potentially explain the somewhat inconclusive existing empirical evidence on the relation between market size and markups; see e.g. Badinger (2007) and Chen et al. (2009). Our results show that the anti-competitive outcomes of an increase in market size sometimes found in these studies may be related to the substantial differences in the distribution of productivities across industries; see Rossi-Hansberg and Wright (2007). This relates to the model of Zhelobodko et al. (2012) in which anti-competitive outcomes may result from the characteristics of consumer preferences rather than collusion.

3.8 Multidimensional Firm Heterogeneity

Although a pattern of firm sorting based solely on productivity is convenient for characterising an equilibrium, it has some undesirable features. First, the strict relationship between productivity and the level of any activity seems unrealistic. This is especially evident when the number of activities is large. Second, this very relationship introduces a gap between the range of available decisions in S and the range of decisions observed in equilibrium.³⁴ Given the wide variety of firm decisions seen in reality, such a limitation on the observable decisions is undesirable.

Therefore, let us consider introducing a vector of firm-specific characteristics other than productivity, $\gamma \in \mathbb{R}^k$, which is a realisation of the random variable Γ , which is distributed independently from θ with c.d.f. G . Firms are now characterised by the pair of characteristics (θ, γ) which are realised simultaneously upon entry. The distinction between productivity and other

condition of a decreasing/nondecreasing hazard rate of the distribution of log-productivity is different from theirs (log-concavity/convexity of this distribution).

³⁴For example, if we have two binary activities, then we can have four possible decisions, $S = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$. However, if $x' = (1, 0)$ is the optimal decision for one firm, then $x'' = (0, 1)$ cannot be the optimal decision for some other firm since this would contradict Lemma 1.

firm characteristics, γ , is made since the assumptions we have made with respect to θ are not imposed on γ . Allowing for additional sources of firm heterogeneity through γ alleviates the two issues mentioned above. While the strict sorting pattern based on productivity holds for a given realisation of γ , it does not necessarily hold across firms with different characteristics, γ . If we consider all firms at once, this can break the strict relationship between the level of a given activity and productivity and thereby increase the number of observable decisions. For examples where multidimensional firm heterogeneity has this purpose, see Eaton et al. (2011), Amiti and Davis (2012), and Hallak and Sivadasan (2013). An implicit assumption in the result below is that the direct effect on s_a is nonpositive for all γ such that the total direct selection effect (across gammas) remains nonpositive.

Proposition 4. *For any distribution G , Propositions 1 and 2 still hold when including multidimensional firm heterogeneity through γ . Proposition 3 also continues to hold for any distribution G if θ_a is independent of γ .*

That Proposition 4 does not require conditions on the distribution G is comforting to the extent that many potential sources of firm heterogeneity are not easily observable. The proposition is proven in Appendix E.³⁵

4 An Application

The present section shows how our firm- and industry-level results can be applied to combine and extend existing models in a way that gives rise to new insights.³⁶

4.1 Trade Liberalisation and Vertical Integration

Bache and Laugesen (2014) show how our results can be used to extend and modify existing models to generate new insights. Specifically, they show that

³⁵It also follows from Appendix E that F and G being independent is unnecessary. It suffices that the needed restrictions on the marginal distribution of θ hold for all γ . Hence, the shape and scale parameters of a Pareto distribution of θ could potentially depend on γ implying that the entire distribution of productivities could look different than the Pareto.

³⁶Our results can also be applied to a number of other models from the trade literature covering a broad range of topics which we do not treat here. Among these are Helpman et al. (2004), Arkolakis (2010), Amiti and Davis (2012), Bustos (2011), and Caliendo and Rossi-Hansberg (2012).

with a slight simplification of the basic setup of Antràs and Helpman (2004), the two activities faced by firms in this model—offshoring and vertical integration of intermediate-input production—are complementary. This is achieved by focusing on the special case where the cost of contractual breach is not affected by offshoring and assuming that fixed costs are linear in the activities. To see how the activities become complementary, note that offshoring implies a reduction of variable costs while vertical integration scales up variable profits (for the relevant firms) due to an improvement of incentives in the producer-supplier relationship. When offshoring does not affect the improvement in the producer-supplier relationship provided by vertical integration, these two activities are complementary in variable profits since scaling up variable profits is worth more when the marginal cost is lower (variable profits are larger to begin with). Assuming that fixed costs are linear (or submodular) in the activities simply ensures that this complementarity in variable profits manifests itself unambiguously in total profits.

Having established this basic complementarity, the model is extended to include within-industry firm heterogeneity with respect to headquarter intensity (as well as productivity) and exporting of final goods as an additional activity. Importantly, exporting is complementary to both offshoring and vertical integration.³⁷ With the results of the present paper in hand, the complementarities inherent in the model imply that these extensions pose no problem for deriving clear comparative statics for the industry composition following liberalisations of either final- or intermediate-goods trade. An interesting implication of the model is that both types of trade liberalisation increase the share of firms that vertically integrate. This is a direct contradiction of Antràs and Helpman (2004) who find that liberalising intermediate-goods trade leads to a reduction in the share of firms that vertically integrate. The divergence in results is a consequence of the desire of Antràs and Helpman (2004) to generate a rich sorting pattern, i.e., observe all combinations of offshoring and vertical integration, based on heterogeneity in productivity alone. This means that they have to make sure that offshoring and vertical integration are not complementary, c.f. Section 3.8, which requires a very specific fixed-cost structure. Bache and Laugesen (2014) show how new insights arise by using headquarter intensity as an additional source

³⁷ Access to another market via exporting increases the production volume, thereby implying that both a lower marginal cost (offshoring) and scaling up variable profits (vertical integration) are more attractive.

of firm-level heterogeneity to obtain a rich sorting pattern and by letting the complementarity between offshoring and vertical integration play out. Further, when these two activities are complementary, the model can easily be extended to include other activities that are complementary to the existing, while maintaining clear comparative statics as exemplified by adding exporting to the model. Finally, as a bonus, the sorting pattern generated by letting firms be heterogeneous with respect to both headquarter intensity and productivity is broadly in line with recent empirical evidence presented by Corcos et al. (2013).

5 Concluding Remarks

One main finding is that firm-level complementarities may manifest themselves much more clearly in the industry composition than in the behaviour of individual firms. Despite the ambiguities at the firm level, we show that one may very well observe that the equilibrium distributions of the activities unambiguously shift towards higher values. Key to this result, which is particularly relevant for international trade and related fields of study, is the observation that these distributions depend not only on the activity levels undertaken by firms conditional on being active but also on the selection of active firms. These results are important for several reasons. First, they provide general insights on the implications of firm-level complementarities in models of monopolistic competition; a workhorse market structure in many strands of the economics literature. Second, our results provide strong, novel, and testable predictions—especially at the industry level—for a large number of recent trade models. We believe that it will be both useful and interesting to confront these predictions with data. On the one hand, such empirical investigations can shed light on the appropriateness of commonly-used functional form assumptions. On the other hand, this approach is likely to complement firm-level and structural estimations. We leave this task for future research. Third, we provide a flexible tool for modelling explanations of shifts in the industry composition based on complementarities at the firm level. Our analysis clearly shows how incentives to undertake one activity at the firm level may induce unambiguous shifts in the equilibrium distributions of this and other activities at the industry level. Fourth and finally, we have illustrated another context in which monotonicity theorems are a powerful mathematical tool for conducting economic analysis.

A Mathematical Appendix

Let $X \subseteq \mathbb{R}^n$ and $T \subseteq \mathbb{R}^m$ be partially ordered sets with the component-wise order.³⁸ For two vectors, $x', x'' \in \mathbb{R}^n$, we let $x' \vee x''$ denote the component-wise maximum and $x' \wedge x''$ denote the component-wise minimum.³⁹ The set X is a lattice if for all $x', x'' \in X$, $x' \vee x'' \in X$ and $x' \wedge x'' \in X$. The set $S \subseteq X$ is a sublattice of X if S is a lattice itself. For two sets, $S', S'' \subseteq \mathbb{R}^n$, we say that S'' is higher than S' and write $S' \leq_s S''$ if for all $x' \in S'$ and all $x'' \in S''$, $x' \vee x'' \in S''$ and $x' \wedge x'' \in S'$. If a set becomes higher, then we say that the set is increasing.

Let X be a lattice. The function $h : X \times T \rightarrow \mathbb{R}$ is supermodular in x on X for each $t \in T$ if for all $x', x'' \in X$ and $t \in T$,

$$h(x', t) + h(x'', t) \leq h(x' \wedge x'', t) + h(x' \vee x'', t). \quad (8)$$

Supermodularity of h in x implies that the return from increasing several elements of x together is larger than the combined return from increasing the elements separately.⁴⁰ This follows from the fact that a higher value of one subset of the elements in x increases the value of increasing other subsets of elements. Supermodularity thus implies that the elements of the vector x are (Edgeworth) complements. If h is smooth, supermodularity is equivalent with $\partial^2 h / \partial x_i \partial x_j \geq 0$ for all i, j where $i \neq j$. By (8), it follows that any function h is trivially supermodular in x when x is a single real variable. The function $h(x, t)$ has increasing differences in (x, t) if for $x' \leq x''$, $h(x'', t) - h(x', t)$ is monotone nondecreasing in t . Increasing differences mean that increasing t raises the return from increasing x and vice versa. If h is smooth, increasing differences are equivalent with $\partial^2 h / \partial x_i \partial t_j \geq 0$ for all i, j . The following monotonicity theorem is due to Topkis (1978).

Theorem 1. *Let $X \subseteq \mathbb{R}^n$ be a lattice, $T \subseteq \mathbb{R}^m$ be a partially ordered set, S be a sublattice of X , and $h : X \times T \rightarrow \mathbb{R}$. If $h(x, t)$ is supermodular in x on X for each $t \in T$ and has increasing differences in (x, t) on $X \times T$, then $\arg \max_{x \in S} h(x, t)$ is monotone nondecreasing in (t, S) .*

³⁸For $x' = (x'_1, \dots, x'_n) \in \mathbb{R}^n$ and $x'' = (x''_1, \dots, x''_n) \in \mathbb{R}^n$, $x' \leq x''$ if $x'_i \leq x''_i$ for $i = 1, \dots, n$ and $x' < x''$ if $x' \leq x''$ and $x' \neq x''$.

³⁹That is, $x' \vee x'' = (\max\{x'_1, x''_1\}, \dots, \max\{x'_n, x''_n\})$ and $x' \wedge x'' = (\min\{x'_1, x''_1\}, \dots, \min\{x'_n, x''_n\})$.

⁴⁰To see this, rewrite (8) into $[h(x', t) - h(x' \wedge x'', t)] + [h(x'', t) - h(x' \wedge x'', t)] \leq h(x' \vee x'', t) - h(x' \wedge x'', t)$.

If the set of maximisers only contains one element, this unique maximiser is nondecreasing in (t, S) . In the remainder of this paper, we restrict attention to cases where the set of maximisers is a nonempty and complete sublattice.⁴¹ This implies that the set of maximisers has greatest and least elements, and Theorem 1 implies that these greatest and least elements are nondecreasing functions of (t, S) . We follow the convention of focusing on the greatest element in the set of maximisers, effectively treating this maximiser as unique.⁴² The monotone comparative statics result of Theorem 1 could also be obtained under the weaker assumption that h is quasisupermodular in x and exhibits the single crossing property in (x, t) .⁴³ However, the properties supermodularity and increasing differences more accurately represent the standard notion of complementarity, are well-known, and are easy to characterise for smooth functions. Therefore we use these assumptions throughout the main part of the paper while keeping in mind that all our results will also hold under quasisupermodularity and the single crossing property.

B Quasisupermodularity and Single Crossing

Let X be a lattice and T be a partially ordered set. The real-valued function $h(x, t)$ is quasisupermodular in x on X if for all $x', x'' \in X$, $h(x', t) \geq h(x' \wedge x'', t)$ implies $h(x' \vee x'', t) \geq h(x'', t)$ and $h(x', t) > h(x' \wedge x'', t)$ implies $h(x' \vee x'', t) > h(x'', t)$. Hence, if an increase in a subset of the elements of x raises h at a given level of the remaining elements, exactly the same increase in the same subset of the elements of x will increase h when the remaining elements also increase. In the language of Milgrom and Shannon (1994), quasisupermodularity expresses a weak kind of complementarity between the elements of x . The function $h(x, t)$ satisfies the single crossing property in (x, t) if for $x'' > x'$ and $t'' > t'$, $h(x'', t') > h(x', t')$ implies that $h(x'', t'') > h(x', t'')$ and $h(x'', t') \geq h(x', t')$ implies that $h(x'', t'') \geq h(x', t'')$. Hence, if an increase in x raises h when t is low, exactly the same increase in x will raise h when t is high. One can verify by the relevant definitions that any supermodular function is also quasisupermodular and any function with

⁴¹General sufficient conditions for this are found in Milgrom and Shannon (1994).

⁴²We share this approach with Bagwell and Ramey (1994) and Holmstrom and Milgrom (1994). All results also hold when one focuses on e.g. the least element.

⁴³See Appendix B for formal definitions and a theorem.

increasing differences in (x, t) also satisfies the single crossing property in (x, t) . Let $S \subseteq X$. The following monotonicity theorem is due to Milgrom and Shannon (1994).

Theorem 2. *$\arg \max_{x \in S} h(x, t)$ is monotone nondecreasing in (t, S) if and only if h is quasisupermodular in x on X for each $t \in T$ and satisfies the single crossing property in (x, t) on $X \times T$.*

C Proof of Propositions 2 and 3

Denote by ΔH_i the change in H_i induced by an increase in (β, S) from (β', S') to (β'', S'') . Using $t' = (\beta', S')$ and $t'' = (\beta'', S'')$ as shorthand notation, this change can be decomposed into the two level and two selection effects mentioned in Section 3.2.

$$\begin{aligned} \Delta H_i = & \underbrace{\frac{s_i(x_i; A', t')}{s_a(A', t')} - \frac{s_i(x_i; A', t'')}{s_a(A', t'')}}_{\text{Direct level effect}} + \underbrace{\frac{s_i(x_i; A', t'')}{s_a(A', t'')} - \frac{s_i(x_i; A', t'')}{s_a(A', t'')}}_{\text{Direct selection effect}} \\ & + \underbrace{\frac{s_i(x_i; A', t'')}{s_a(A', t'')} - \frac{s_i(x_i; A'', t'')}{s_a(A', t'')}}_{\text{Indirect level effect}} + \underbrace{\frac{s_i(x_i; A'', t'')}{s_a(A', t'')} - \frac{s_i(x_i; A'', t'')}{s_a(A'', t'')}}_{\text{Indirect selection effect}}, \quad (9) \end{aligned}$$

where $A' = A(\beta', S')$ and $A'' = A(\beta'', S'')$. The total level effect is due to changes in the levels of activity i undertaken by firms conditional on being active. In (9), this is represented by changes in the share of firms undertaking at least a given level of activity i , s_i . The total selection effect is due to changes in the range of active firms which is represented by changes in the share of active firms, s_a . Each of these two total effects has a direct component induced by changes in (β, S) for a given A and an indirect component induced by changes in A .

Note that a nonpositive (total) indirect effect in (9) is equivalent to

$$\frac{s_i(x_i; A', t'')}{s_a(A', t'')} \leq \frac{s_i(x_i; A'', t'')}{s_a(A'', t'')}. \quad (10)$$

When F is C^1 , we have that $1 - F(\theta) = e^{-\int_{\theta_0}^{\theta} \lambda_{\theta}(u) du}$, where λ_{θ} denotes the hazard rate of the distribution of θ . Using this observation, (10) can be expressed as

$$e^{-\int_{\theta_a(A', \beta'', S'')}^{\theta_a(A'', \beta'', S'')} \lambda_{\theta}(u) du} \leq e^{-\int_{\theta_i(x_i; A', \beta'', S'')}^{\theta_i(x_i; A'', \beta'', S'')} \lambda_{\theta}(u) du}.$$

Note that when the profit function depends on A and θ only through $A\theta$, the changes in $\log \theta_i$ and $\log \theta_a$ induced by a change in A are equal. A change of integrand gives us

$$e^{-\int_{\log \theta_a(A', \beta'', S'')}^{\log \theta_a(A'', \beta'', S'')} \lambda_{\log \theta}(u) du} \leq e^{-\int_{\log \theta_i(x_i; A', \beta'', S'')}^{\log \theta_i(x_i; A'', \beta'', S'')} \lambda_{\log \theta}(u) du}, \quad (11)$$

where $\lambda_{\log \theta}$ is the hazard rate of the distribution of log-productivity. Now, since $\theta_i \geq \theta_a$ and since the changes induced by a change in A in $\log \theta_i$ and $\log \theta_a$ are equal and nonnegative if competition is enhanced, the condition (11) is fulfilled if the hazard rate of log-productivity is nonincreasing. On the other hand, since the changes induced by a change in A in $\log \theta_i$ and $\log \theta_a$ are equal and nonpositive if competition is dampened, the condition (11) is fulfilled if the hazard rate of log-productivity is nondecreasing. It also follows that, if the hazard rate of log-productivity is constant, the condition (11) is satisfied both when competition is enhanced and when it is dampened. Finally, the distribution of $\log \theta$ having constant hazard rate is equivalent with θ being Pareto distributed.⁴⁴

D The Dependence of Profits on A and θ

Consider the assumption that profits depend on θ and A only through their product, $A\theta$. First, we emphasise that some degree of separability—where we focus on multiplicative separability—between θ and A in profits is necessary in order to make the two indirect effects on H_i balance and thereby achieve MCS for the industry composition regardless of the change in competition. To see this, note that the requirement for the indirect effects on H_i to balance can be expressed as⁴⁵

$$\int_{\theta_a(A', \beta'', S'')}^{\theta_a(A'', \beta'', S'')} \lambda_{\theta}(u) du = \int_{\theta_i(x_i; A', \beta'', S'')}^{\theta_i(x_i; A'', \beta'', S'')} \lambda_{\theta}(u) du, \quad (12)$$

where λ_{θ} is the hazard rate of the distribution of productivity. This gives us a functional relationship between the distribution (hazard rate) of productivity and the way the decision on activities and the choice to exit depends on A

⁴⁴ $\log \theta$ being distributed with constant hazard rate, $\lambda_{\log \theta}$, implies that $F_{\log \theta}(\log \theta) = 1 - e^{-\lambda_{\log \theta}(\log \theta - \log \theta_0)}$, where $F_{\log \theta}$ denotes the c.d.f. of $\log \theta$. Rearranging gives $F(\theta) = 1 - (\theta_0/\theta)^{\lambda_{\log \theta}}$. Thus $F(\theta)$ is given by the Pareto distribution.

⁴⁵See Appendix C.

and θ . To get a sense of this relationship, consider a change of variable in (12) such that the new integrand, $\varphi = \varphi(\theta)$, is distributed with constant hazard rate, λ_φ .⁴⁶ It follows that

$$\int_{\varphi_a(A',\beta'',S'')}^{\varphi_a(A'',\beta'',S'')} \lambda_\varphi(u) du = \int_{\varphi_i(x_i;A',\beta'',S'')}^{\varphi_i(x_i;A'',\beta'',S'')} \lambda_\varphi(u) du, \quad (13)$$

where $\varphi_a = \varphi(\theta_a)$ and $\varphi_i = \varphi(\theta_i)$. For (13) to hold true, we must have that the changes in φ_a and φ_i are of the exact same size. This condition readily gives you that the decision on the activities and the choice to exit depend on φ and A only through $\varphi + Z(A, \beta, S)$ where Z is an arbitrary function which is nondecreasing in A . This is very closely connected to profits only depending on φ and A through $\varphi + Z(A, \beta, S)$. Thus, to make sure that the total indirect effect of an increase in (β, S) on H_i is zero, additive separability of (transformations of) productivity and the demand level is central. Now, if profits only depend on θ and A through $A\theta$, this can also be expressed as profits only depending on θ and A through $\log \theta + \log A$. Using $\varphi(\theta) = \log \theta$ and $Z(A, \beta, S) = \log A$, the condition (12) is satisfied if $\log \theta$ is distributed with constant hazard rate, i.e., if θ is Pareto distributed.

E Proof of Proposition 4

It is obvious that Proposition 1 is unaffected by the introduction of multi-dimensional firm heterogeneity. Next, note that with multidimensional firm heterogeneity, as introduced in Section 3.8, we can express

$$H_i(x_i; \beta, S) = 1 - \frac{\int s_i(x_i, \gamma; A(\beta, S), \beta, S) dG(\gamma)}{\int s_a(\gamma; A(\beta, S), \beta, S) dG(\gamma)},$$

where θ_i and θ_a now (possibly) depend on γ , wherefore s_i and s_a do so as well. Start by noting that the total direct selection effect remains nonpositive (see main text). This is also the case for the total direct level effect since the direct level effect for a given γ is nonpositive. Repeating the steps in Appendix C, the condition for a nonpositive (total) indirect effect on H_i becomes

$$\int \omega_a e^{-\int_{\log \theta_a(\gamma; A', \beta'', S'')}^{\log \theta_a(\gamma; A'', \beta'', S'')} \lambda_{\log \theta}(u) du} dG(\gamma) \leq \int \omega_i e^{-\int_{\log \theta_i(x_i, \gamma; A', \beta'', S'')}^{\log \theta_i(x_i, \gamma; A'', \beta'', S'')} \lambda_{\log \theta}(u) du} dG(\gamma), \quad (14)$$

⁴⁶Letting $\varphi(\theta) = -\log(1 - F(\theta))/\lambda_\varphi$ ensures this.

where we have defined the weights $\omega_a \equiv s_a(\gamma; A', \beta'', S'') / \int s_a(\gamma; A', \beta'', S'') dG(\gamma)$ and $\omega_i \equiv s_i(x_i, \gamma; A', \beta'', S'') / \int s_i(x_i, \gamma; A', \beta'', S'') dG(\gamma)$ which both integrate to one. When the hazard rate of log-productivity is constant, it is obvious that this condition is still satisfied both when competition is enhanced and when it is dampened since the changes in $\log \theta_i$ and $\log \theta_a$ are still equal. Hence, Proposition 2 still holds.

When θ_a is independent of γ , the condition (14) simplifies to

$$e^{-\int_{\log \theta_a(A', \beta'', S'')}^{\log \theta_a(A'', \beta'', S'')} \lambda_{\log \theta}(u) du} \leq \int \omega_i e^{-\int_{\log \theta_i(x_i, \gamma; A', \beta'', S'')}^{\log \theta_i(x_i, \gamma; A'', \beta'', S'')} \lambda_{\log \theta}(u) du} dG(\gamma),$$

and Proposition 3 holds for the same reasons it did without multidimensional firm heterogeneity.

References

- Acemoglu, D., Jensen, M.K., 2013. Aggregate Comparative Statics. *Games and Economic Behaviour* 81, 27–49.
- Amir, R., 2005. Supermodularity and Complementarity in Economics: An Elementary Survey. *Southern Economic Journal* 71(3), 636–660.
- Amiti, M., Davis, D.R., 2012. Trade, Firms, and Wages: Theory and Evidence. *Review of Economic Studies* 79, 1–36.
- Antràs, P., Helpman, E., 2004. Global Sourcing. *Journal of Political Economy* 112(3), 552–580.
- Antràs, P., Fort, T.C., Tintelnot, F., 2014. The Margins of Global Sourcing: Theory and Evidence from U.S. Firms. Working Paper.
- Arkolakis, C., 2010. Market Penetration Costs and the New Consumers Margin in International Trade. *Journal of Political Economy* 118(6), 1151–1199.
- Arkolakis, C., Costinot, A., Donaldson, D., Rodríguez-Clare, A., 2012. The Elusive Pro-Competitive Effects of Trade. Working Paper.
- Athey, S., 2002. Monotone Comparative Statics under Uncertainty. *The Quarterly Journal of Economics* 117(1), 187–223.

- Athey, S., Schmutzler, A., 1995. Product and Process Flexibility in an Innovative Environment. *RAND Journal of Economics* 26(4), 557–574.
- Axtell, R.L., 2001. Zipf Distribution of U.S. Firm Sizes. *Science* 293(5536), 1818–1820.
- Bache, P.A., Laugesen, A., 2014. Trade Liberalisation and Vertical Integration. Working Paper.
- Badinger, H., 2007. Has the EU’s Single Market Programme Fostered Competition? Testing for a Decrease in Mark-up Ratios in EU Industries. *Oxford Bulletin of Economics and Statistics* 69(4), 497–519.
- Bagwell, K., Ramey, G., 1994. Coordination Economies, Advertising, and Search Behavior in Retail Markets. *American Economic Review* 84(3), 408–517.
- Bellone, F., Musso, P., Nesta, L., Warzynski, F., 2014. International Trade and Firm-level Markups when Location and Quality Matter. *Journal of Economic Geography* , 1–25.
- Bernard, A.B., Eaton, J., Jensen, J.B., Kortum, S., 2003. Plants and Productivity in International Trade. *American Economic Review* 93(4), 1268–1290.
- Bernard, A.B., Redding, S.J., Schott, P.K., 2011. Multiproduct Firms and Trade Liberalization. *The Quarterly Journal of Economics* 126(3), 1271–1318.
- Bustos, P., 2011. Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms. *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(3), 1393–1467.
- Chen, N., Imbs, J., Scott, A., 2009. The Dynamics of Trade and Competition. *Journal of International Economics* 77(1), 50–62.
- Combes, P.P., Duranton, G., Gobillon, L., Puga, D., Roux, S., 2012. The Productivity Advantages of Large Cities: Distinguishing Agglomeration from Firm Selection. *Econometrica* 80(6), 2543–2594.

- Corcos, G., Irac, D., Mion, G., Verdier, T., 2013. The Determinants of Intrafirm Trade: Evidence from French Firms. *Review of Economics and Statistics* 95(3), 825–838.
- Costinot, A., 2007. Heterogeneity and Trade. Working Paper.
- Costinot, A., 2009. An Elementary Theory of Comparative Advantage. *Econometrica* 77(4), 1165–1192.
- Davis, D.R., Harrigan, J., 2011. Good Jobs, Bad Jobs, and Trade Liberalization. *Journal of International Economics* 84, 26–36.
- Eaton, J., Kortum, S., 2002. Technology, Geography, and Trade. *Econometrica* 70(5), 1771–1780.
- Eaton, J., Kortum, S., Kramarz, F., 2011. An Anatomy of International Trade: Evidence from French Firms. *Econometrica* 79(5), 1453–1498.
- Eeckhout, J., 2004. Gibrat’s Law for (all) Cities. *American Economic Review* 94(5), 1429–1451.
- Egger, H., Koch, M., 2013. Trade and the Firm-Internal Allocation of Workers to Tasks. Working Paper.
- Forslid, R., Okubo, T., Ulltveit-Moe, K.H., 2014. Why are Firms that Export Cleaner? *International Trade and CO2 Emmisions*. Working Paper.
- Hallak, J.C., Sivadasan, J., 2013. Product and Process Productivity: Implications for Quality Choice and Conditional Exporter Premia. *Journal of International Economics* 91, 53–67.
- Head, K., Mayer, T., Thoenig, M., 2014. Welfare and Trade without Pareto. *American Economic Review* 104(5), 310–316.
- Helpman, E., Itskhoki, O., 2010. Labour Market Rigidities, Trade and Unemployment. *Review of Economic Studies* 77, 1100–1137.
- Helpman, E., Itskhoki, O., Redding, S.J., 2010. Inequality and Unemployment in a Global Economy. *Econometrica* 78(4), 1239–1283.
- Helpman, E., Melitz, M.J., Yeaple, S.R., 2004. Export Versus FDI with Heterogeneous Firms. *American Economic Review* 94(1), 300–316.

- Holmstrom, B., Milgrom, P., 1994. The Firm as an Incentive System. *American Economic Review* 84(4), 972–991.
- Hopenhayn, H., 1992. Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica* 60, 1127–1150.
- Ijiri, Y., Simon, H.A., 1974. Interpretations of Departures from the Pareto Curve Firm-Size Distributions. *Journal of Political Economy* 82(2), 315–331.
- Jensen, M.K., 2014. Distributional Comparative Statics. Working Paper.
- Jeremy, B., Geanakoplos, J., Klemperer, P., 1985. Multimarket Oligopoly: Strategic Substitutes and Complements. *Journal of Political Economy* 93, 488–511.
- Krugman, P., 1979. Increasing Returns, Monopolistic Competition, and International Trade. *Journal of International Economics* 9(4), 469–479.
- Luttmer, E.G.J., 2007. Selection, Growth, and the Size Distribution of Firms. *The Quarterly Journal of Economics* 122, 1103–1144.
- Mayer, T., Melitz, M., Ottaviano, G.I., 2014. Market Size, Competition, and the Product Mix of Exporters. *American Economic Review* 104(2), 495–536.
- Melitz, M.J., 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica* 71(6), 1695–1725.
- Melitz, M.J., Ottaviano, G.I.P., 2008. Market Size, Trade, and Productivity. *Review of Economic Studies* 75, 295–316.
- Melitz, M.J., Redding, S.J., 2014. Heterogeneous Firms and Trade, in: *Handbook of International Economics*. Elsevier.
- Milgrom, P., Qian, Y., Roberts, J., 1991. Complementarities, Momentum, and the Evolution of Modern Manufacturing. *AEA Papers and Proceedings* 81(2), 84–88.
- Milgrom, P., Roberts, J., 1990a. The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *American Economic Review* 80(3), 511–528.

- Milgrom, P., Roberts, J., 1990b. Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarity. *Econometrica* 58(6), 1255–1277.
- Milgrom, P., Roberts, J., 1995. Complementarities and Fit Strategy, Structure, and Organizational Change in Manufacturing. *Journal of Accounting and Economics* 19, 179–208.
- Milgrom, P., Roberts, J., 1996. The LeChatelier Principle. *American Economic Review* 86(1), 173–179.
- Milgrom, P., Shannon, C., 1994. Monotone Comparative Statics. *Econometrica* 62(1), 157–180.
- Mrazova, M., Neary, J.P., 2013. Selection Effects with Heterogeneous Firms. Working Paper.
- Ottaviano, G., Tabuchi, T., Thisse, J.F., 2002. Agglomeration and Trade Revisited. *International Economic Review* 43(2), 409–436.
- Rossi-Hansberg, E., Wright, M.L.J., 2007. Establishment Size Dynamics in the Aggregate Economy. *American Economic Review* 97(5), 1639–1666.
- Topkis, D.M., 1978. Minimizing a Supermodular Function on a Lattice. *Operations Research* 26, 305–321.
- Topkis, D.M., 1979. Equilibrium Points in Nonzero-sum n -person Supermodular Games. *SIAM Journal of Control and Optimization* 17, 773–787.
- Topkis, D.M., 1995. Comparative Statics of the Firm. *Journal of Economic Theory* 67(2), 370–401.
- Vives, X., 1990. Nash Equilibrium with Strategic Complementarities. *Journal of Mathematical Economics* 19(3), 305–321.
- Zhelobodko, E., Kokovin, S., Parenti, M., Thisse, J., 2012. Monopolistic Competition: Beyond the C.E.S. *Econometrica* 80(6), 2765–2784.