



DEPARTMENT OF ECONOMICS
DISCUSSION PAPER SERIES

SELECTION EFFECTS WITH HETEROGENEOUS FIRMS

Monika Mrazova and J. Peter Neary

Number 588
December 2011

Manor Road Building, Oxford OX1 3UQ

SELECTION EFFECTS WITH HETEROGENEOUS FIRMS*

Monika Mrázová[†]

University of Surrey
and CEP, LSE

J. Peter Neary[‡]

University of Oxford
and CEPR

December 24, 2011

Abstract

We provide a general characterization of which firms will select alternative ways of serving a market. If and only if firms' maximum profits are supermodular in production and market-access costs, more efficient firms will select into the activity with lower market-access costs. Our result applies in a range of models and under a variety of assumptions about market structure. We show that supermodularity holds in many cases but not in all. Exceptions include FDI (both horizontal and vertical) when demands are “sub-convex” (i.e., less convex than CES), fixed costs that vary with access mode, and R&D with threshold effects.

Keywords: Foreign Direct Investment (FDI); Heterogeneous Firms; Proximity-Concentration Trade-Off; R&D with Threshold Effects; Super- and Sub-Convexity; Supermodularity.

JEL Classification: F23, F15, F12

*We are grateful to Arnaud Costinot for extensive discussions, and to Paola Conconi, Jonathan Dingle, Carsten Eckel, Peter Egger, Gene Grossman, Dermot Leahy, Marc Melitz, Giordano Mion, Lindsay Oldenski, Emanuel Ornelas, Gianmarco Ottaviano, Adrian Wood, Kresimir Zigic, and participants at various seminars and conferences for helpful comments. Monika Mrázová wishes to thank the ESRC [grant number PTA-026-27-2479].

[†]School of Economics, University of Surrey, Guildford, Surrey GU2 7XH, UK; e-mail: m.mrazova@surrey.ac.uk; URL: <http://www.monikamrazova.com>.

[‡]Department of Economics, University of Oxford, Manor Road, Oxford OX1 3UQ, UK; e-mail: peter.neary@economics.ox.ac.uk; URL: <http://www.economics.ox.ac.uk/members/peter.neary/neary.htm>.

1 Introduction

Why do different firms choose to serve particular markets in different ways? Ten years ago, economists had little theory to guide them in thinking about such questions, though a growing body of empirical work had already documented systematic patterns in firm-level data that were unexplained by traditional theory. In the intervening decade, a new and exciting body of theoretical work has emerged which has placed these empirical findings in context and inspired further extensions and elaborations. The starting point of this recent literature is the explicit recognition that firms differ in one or more underlying attribute, typically identified with their productivity; and its central prediction is that more productive firms select into activities with higher fixed costs but lower variable costs. The *locus classicus* for this pattern of behavior is Melitz (2003), who showed that more efficient firms select into exporting, whereas less efficient ones serve the home market only.¹ Subsequent work in the same vein has shown that more efficient firms select into many different activities, such as producing in-house rather than outsourcing, as in Antràs and Helpman (2004); serving foreign markets via foreign direct investment (FDI) rather than exports, as in Helpman, Melitz, and Yeaple (2004); paying higher wages as in Egger and Kreickemeier (2009) and Helpman, Itskhoki, and Redding (2010); and producing with more skill-intensive techniques as in Bustos (2011). Exploring the implications of firm heterogeneity has already had a profound effect on the study of international trade, and is increasingly being extended to other fields, for example by Ghironi and Melitz (2005) to international macroeconomics, by Davies and Eckel (2010) to international tax competition, and by Forslid, Okubo, and Ulltveit-Moe (2011) to environmental economics.

This literature on heterogeneous firms prompts a number of observations. First, international trade is not the only field in economics where it has been noted that a firm's superiority in one dimension may enhance its performance in others. The same

¹A related result of this recent literature, also due to Melitz (2003), is a new source of gains from trade: trade liberalization encourages exit by less productive firms and entry by more productive ones, and so, even when the productivity of each individual firm is unchanged, aggregate productivity rises. However, recent work by Arkolakis, Costinot, and Rodríguez-Clare (2009) suggests that this effect operates in a similar fashion to the gains from trade in traditional models with homogeneous firms.

idea, though expressed in very different ways, can be found in Milgrom and Roberts (1990), who argued that such a complementarity or “supermodularity” between different aspects of firm performance is typical of modern manufacturing. They also advocated using the mathematical tools of robust comparative statics to examine the responses of such firms to exogenous shocks, especially in contexts where variables may change by discrete amounts. This suggests that it may be worth exploring possible links between these two literatures, and possible payoffs to adapting the tools of robust comparative statics to better understand the behavior of heterogeneous firms.

Second, the question immediately arises whether the results derived to date in the literature on heterogeneous firms and trade are robust. One dimension of robustness is that of functional form. All the papers cited above assume that consumers have Dixit-Stiglitz or constant-elasticity-of-substitution (CES) preferences, and all but Melitz (2003) assume that firm productivities follow a Pareto distribution. These assumptions have been relaxed in some papers; for example, Melitz and Ottaviano (2008) show that more efficient firms also select into exports when preferences are quadratic rather than CES. However, with existing techniques each small change in assumptions requires that the model be solved again in full, and, as a result, the question of robustness with respect to functional form has been relatively little explored. A different dimension of robustness is symmetry: existing studies typically assume that countries are identical, both in size and in the distribution of firm productivities. Does this matter for the results? Yet another dimension of robustness is market structure. All the literature on heterogeneous firms to date assumes that the industry is monopolistically competitive, so firms produce a unique product but are infinitesimal in their market. However, if successful firms are indeed large in every dimension, then monopolistic competition may not be an appropriate assumption about market structure. At least in some markets, it may be more plausible to allow for the emergence of a small number of large firms, competing strategically against each other, and possibly coexisting with a “monopolistically competitive fringe” as in Neary (2010). Clearly it would be desirable to know if the selection effects that have been derived assuming monopolistic competition are also likely to hold in oligopolistic

markets. Finally, turning robustness on its head, we can ask whether the fact that more efficient firms engage in more activities is a universal tendency. Should we always expect more productive firms to engage in more and more complex activities? Or are there interesting counter-examples?

In this paper we seek to illuminate these issues both substantively and technically. Our main substantive contribution is a general result on firm selection which fully characterizes the conditions under which what we call the “conventional sorting” pattern occurs: more efficient firms select into activities with lower marginal costs. We first prove this result in a simple though canonical context: that of a single monopoly firm choosing between serving a foreign market by either exports or horizontal FDI. We then show that, with appropriate qualifications, the result extends to a wide variety of market structures, including both monopolistic competition and oligopoly; and to a wide variety of firm decisions, including vertical FDI, in-house production versus outsourcing, and choice of technique. In all cases, the key consideration is how a firm’s own marginal cost of production interacts with the marginal cost of serving the market under different access modes. Our result reveals the unifying structure underlying a wide range of results in the literature, and also shows how they can easily be generalized in new and important ways.

From a technical point of view, the paper contributes to the small but growing literature which uses the techniques of monotone comparative statics, and in particular the concept of supermodularity, to illuminate issues in international trade.² Supermodularity arises very naturally in our context. Our interest is in comparing firms whose production costs differ by a finite amount, and in particular in comparing their behavior under different modes of serving a market, whose marginal costs also differ by a finite amount. Supermodularity imposes a natural restriction on the finite “difference-in-differences” of

²For more technical details on the application of monotone comparative statics to economics, see Milgrom and Roberts (1990), Milgrom and Shannon (1994), and Athey (2002). Other applications of supermodularity to international trade include Grossman and Maggi (2000), Costinot (2009), and Costinot and Vogel (2010), who use it to study problems of matching between different types of workers or between workers and sectors; Limão (2005), who considers links between trade and environmental agreements; and Costinot (2007), who considers firm selection effects, though in a very different way to the present paper.

the firm's profit function which we need to sign in order to make this comparison. As we show, the profit function exhibits supermodularity under a wide range of assumptions, which allows us to generalize existing results and derive new ones with remarkably few restrictions on technology, tastes, or market structure.

The plan of the paper is as follows. Sections 2 and 3 focus on a single monopoly firm which faces the decision on how to serve a foreign market, trading off the proximity benefits of foreign direct investment against the concentration advantage of producing at home and exporting. Section 2 introduces the setting and explains the restrictions implied by supermodularity. Section 3 formalizes the gains from tariff-jumping and derives our main result on how firms of different productivities will select into one or other mode of serving the foreign market. The remainder of the paper shows that our approach applies in a wide range of contexts, both old (including some of the most widely-used models in international trade), and new. Sections 4 and 5 look at alternative market structures, considering monopolistic competition and oligopoly respectively. Section 6 turns to explore firm choices other than that between exports and FDI. It reviews a range of other applications, and considers the implications of heterogeneous and endogenous fixed costs of production. The overall message of the paper is that supermodularity holds in many cases but is not inevitable. Among the specific examples we give where supermodularity may be violated, and so the conventional assignment of firms to different modes of accessing foreign markets may be reversed, are FDI (both horizontal and vertical) with non-CES preferences, fixed costs that vary with access mode, and R&D with threshold effects.

2 Operating Profits and Supermodularity

In this section we consider a firm located in one country which contemplates serving consumers located in a foreign country. The maximum operating profits the firm can earn in the foreign country equal $\pi(t, c)$, where t is the access cost (tariffs and transport costs) it faces and c is an exogenous cost parameter, which we can think of as an inverse

measure of productivity. The parameter c equals the firm's marginal production cost in many applications, though not in all: we will see exceptions in Example 2 and Section 6 below. We assume that π is non-increasing (though not necessarily continuous) in both t and c . Profits also depend on the firm's choice variables and on other exogenous variables. However, the former have been chosen optimally and so are subsumed into the π function, while the latter are suppressed for convenience; we give some examples of each below.

We define Δ_c as the finite difference between the values of a function evaluated at two different values of c , c_1 and c_2 , with the convention that c_1 is greater than or equal to c_2 . Applying this to the operating profit function π gives:

$$\Delta_c \pi(t, c) \equiv \pi(t, c_1) - \pi(t, c_2) \quad \text{when} \quad c_1 \geq c_2 \quad (1)$$

So, $\Delta_c \pi(t, c)$ is the profit loss of a higher-cost relative to a lower-cost firm and is always non-positive. Note that, when $\pi(t, c)$ is differentiable in c , $\frac{\Delta_c \pi(t, c)}{c_1 - c_2}$ reduces to the partial derivative π_c as c_1 approaches c_2 .³

We can now define what we mean by supermodularity in the context of our paper:⁴

Definition 1. The function $\pi(t, c)$ is supermodular in t and c if and only if:

$$\Delta_c \pi(t_1, c) \geq \Delta_c \pi(t_2, c) \quad \text{when} \quad t_1 \geq t_2. \quad (2)$$

Intuitively, supermodularity of π means that a higher tariff reduces in absolute value the cost disadvantage of a higher-cost firm. Putting this differently, the profit function exhibits the “Matthew Effect”: “to those who have, more shall be given”. Rewriting the definition we can see that supermodularity is equivalent to:

$$\pi(t_2, c_2) - \pi(t_1, c_2) \geq \pi(t_2, c_1) - \pi(t_1, c_1) \geq 0 \quad \text{when} \quad t_2 \leq t_1 \quad \text{and} \quad c_2 \leq c_1 \quad (3)$$

³We use subscripts of functions to denote partial derivatives: e.g., $\pi_c \equiv \partial \pi / \partial c$ and $\pi_{tc} \equiv \partial^2 \pi / \partial t \partial c$.

⁴More generally, following Milgrom and Roberts (1990) and Athey (2002), supermodularity can be defined in terms of vector-valued arguments: π is supermodular in a vector-valued argument when $\pi(x \vee y) + \pi(x \wedge y) \geq \pi(x) + \pi(y)$, where $x \vee y \equiv \inf \{z \mid z \geq x, z \geq y\}$ and $x \wedge y \equiv \sup \{z \mid z \leq x, z \leq y\}$. This is equivalent to the definition in the text when we set: $x = \{c_2, t_1\}$ and $y = \{c_1, t_2\}$.

Thus, when supermodularity holds, a lower tariff is of more benefit to a more productive firm. This might seem like the natural outcome, since a lower tariff contributes more to profits the more a firm sells, and we might expect a more productive firm to sell more. As we will see, this is often the case, but there are important counter-examples. When the first inequality in (2) is reversed, we say that the function is submodular.

Note that, when $\pi(t, c)$ is differentiable in t and c , supermodularity of π implies that the second derivative π_{tc} is positive as t_1 approaches t_2 and c_1 approaches c_2 . In the differentiable case, supermodularity is analogous to Hicksian complementarity in consumer theory or strategic complementarity in game theory.

Example 1. A simple case which helps to fix ideas is that of a single-product monopoly firm with constant marginal cost and specific tariffs. Let $p(x)$ denote the inverse demand function which the firm faces, where p and x denote its price and sales respectively. Its operating profits therefore equal:

$$\pi(t, c) \equiv \underset{x}{Max} [\{p(x) - c - t\}x] \quad (4)$$

It is easy to check that the profit function is supermodular in t and c in this case.⁵ Intuitively, a firm with higher production costs c has lower sales; hence its profits are reduced less by a rise in the tariff.

Example 1 exhibits two key features: π is continuous in trade and production costs, and it depends only on their sum. If both these conditions hold, then supermodularity in t and c is equivalent to *convexity* of π in both t and c : if $\pi(t, c) = \pi(t + c)$ and π is differentiable, then $\pi_{tc} = \pi_{cc}$. Our next example is a simple case where one of these conditions does not hold and as a result the profit function may not exhibit supermodularity.

Example 2. Consider next the same example as above except that marginal cost varies

⁵It follows from the envelope theorem that the first derivative of π with respect to t is minus the initial level of sales: $\pi_t = -x(t, c)$. Hence the second cross-partial derivative of profits is minus the partial derivative of sales with respect to c : $\pi_{tc} = -x_c > 0$. To establish the sign of this term, differentiate the first-order condition $p - c - t + xp' = 0$ to get: $x_c = -H^{-1}$. The expression $H \equiv -(2p' + xp'')$ must be positive from the firm's second-order condition. Hence we have that $\pi_{tc} = -x_c = H^{-1} > 0$, and so π in (4) is supermodular in t and c .

with output.⁶ Assume the firm's problem is as follows:

$$\pi(t, c) \equiv \underset{x}{Max} [\{p(x) - C(c, x) - t\}x] \quad (5)$$

Here c is not equal to marginal cost, but rather it is a parameter representing the firm's inverse productivity. The new expression $C(c, x)$ is the firm's average variable cost: it depends positively on c and also varies with output x .⁷ (As we will see, the sign of the first derivative C_x , which indicates whether average variable cost rises or falls with output, is irrelevant to whether supermodularity obtains.) Now the second cross-derivative of the profit function equals the following:⁸

$$\pi_{tc} = H^{-1} (C_c + xC_{xc}) \quad (6)$$

The term H is positive from the second-order condition for profit maximization, and C_c in the expression in parentheses is positive as already noted, both of which work in favor of supermodularity. However, these terms could be offset, and the profit function could be submodular, if C is *submodular* in its arguments $\{x, c\}$ so the term C_{xc} is negative; that is, if the cost of production falls faster (or rises more slowly) with output for a firm with higher c (i.e., a less productive firm). Figure 1 illustrates this possibility. Firm 1 is less productive than firm 2 overall, but it is *relatively* more productive at higher levels of output. As a result its marginal cost curve MC_1 lies below that of firm 2 and so it has *lower* marginal cost and (facing the same marginal revenue curve) *higher* output. The profit function in this case is therefore submodular rather than supermodular.

The configuration shown in Figure 1, though not pathological, is somewhat contrived and of limited empirical relevance. In general, supermodularity will hold as long as static differences in efficiency between firms work in the same direction on average and at the margin, which seems the natural case. In later sections we will consider more plausible

⁶We are grateful to Dermot Leahy for suggesting this example.

⁷Total variable cost equals $x C(c, x)$, so marginal cost is $C + x C_x$, which also varies with output.

⁸As in Example 1, the envelope theorem holds, $\pi_t = -x(t, c)$, and so π is supermodular in t and c if and only if x is decreasing in c : $\pi_{tc} = -x_c$. Direct calculation yields equation (6) where $H \equiv -[2p' + xp'' - (2C_x + xC_{xx})] > 0$.




Figure 1: An Example of Submodularity

examples of submodularity.

3 Selection into FDI versus Exporting

We return to the general case where $\pi(t, c)$ is unrestricted, and compare the relative profitability of different modes of serving the foreign market. We first restate in our notation the familiar proximity-concentration trade-off, and then derive a general result on which firms will select into exporting and FDI.⁹

Exporting faces a higher access cost, so FDI has the advantage of proximity. However, it foregoes the benefits of concentration. In addition to operating profits, the firm must incur a fixed cost of serving the market, which differs depending on the mode of access. The fixed cost equals f_X if the firm exports and f_F if it engages in FDI and builds a plant in the market in question. We assume that f_F is strictly greater than f_X and that access costs conditional on FDI are zero.

Total profits as a result of exporting are:

$$\Pi^X = \pi(t, c) - f_X \quad (7)$$

⁹Our formalization of the proximity-concentration trade-off follows Neary (2002).

while total profits from locating a plant in the target market are:

$$\Pi^F = \pi(0, c) - f_F \quad (8)$$

We define the *tariff-jumping gain* γ as the difference between these two:¹⁰

$$\gamma(t, c, f) \equiv \pi(0, c) - \pi(t, c) - f = \Pi^F - \Pi^X \quad (9)$$

Here $f \equiv f_F - f_X$ is the excess fixed cost of FDI relative to exporting, which by our earlier assumption is strictly positive. We can now apply the finite difference operator Δ_c to the tariff-jumping gain:¹¹

$$\Delta_c \gamma(t, c, f) = \Delta_c \pi(0, c) - \Delta_c \pi(t, c) \quad (10)$$

Recalling the definition of supermodularity, we can sign this unambiguously, which gives our first result:

Lemma 1. *If and only if the profit function π is supermodular in t and c , $\Delta_c \gamma(t, c, f)$ is negative.*

The economic implications of this are immediate: if and only if π is supermodular in t and c , the tariff-jumping gain is lower for higher-cost firms and higher for more productive ones. Since γ measures the incentive to engage in FDI relative to exporting, we can go further and state the central result of our paper:

Proposition 1. *Assume that some firms serve the foreign market by exports and the remainder by FDI. Then, if and only if the profit function π is supermodular in t and c , higher-cost firms will select into exports, while lower-cost firms will select into FDI, for all admissible fixed costs f .*

The sufficiency part of the proposition follows immediately from Lemma 1. The necessity

¹⁰Strictly γ is a *trade-cost-jumping gain*, but the shorter title is traditional and simpler.

¹¹To avoid confusion, we include f among the arguments of $\Delta_c \gamma(t, c, f)$. However, this finite difference is independent of f , a point which will become crucial below.

part is more subtle and reflects the fact that we require the result to hold for all admissible fixed costs. A formal proof is in the Appendix. Here we give an intuitive account.




Figure 2: The Conventional Sorting can Hold without Supermodularity

The upper quadrant of Figure 2, based on Helpman, Melitz, and Yeaple (2004), gives a hypothetical illustration of total profits under the two modes of market access, as functions of inverse production costs.¹² The lower quadrant shows $\gamma(t, c, f)$, the difference between Π^F and Π^X . In the example shown, this curve is non-monotonic in c , and so the profit function is not supermodular in t and c . Despite this, the conventional sorting holds: there is a unique cost threshold, with all firms that have lower costs engaging in FDI and all those that have higher costs engaging in exports. This shows that supermodularity is not necessary for the conventional sorting to hold for arbitrary fixed costs: the necessary condition, that the two profit curves cross only once, is weaker. However, Figure 3 shows that supermodularity is necessary if the conventional sorting is to hold for *all* fixed costs. If γ is not monotonic in c , there always exists some level of f which leads to a violation of the conventional sorting. In the case shown in Figure 3, the level of f_F , the fixed cost of FDI, is lower than in Figure 2, shifting the $\gamma(t, c, f)$ curve upwards such that the

¹²Helpman, Melitz, and Yeaple (2004) assume CES preferences, in which case the total profit curves are linear in a decreasing transformation of costs; see Section 4 below for further discussion. More generally, the total profit curves must satisfy only two restrictions: (i) both must be upward-sloping, reflecting the assumption that operating profits are non-increasing in c ; and (ii) Π^F must lie everywhere above Π^X when $f_F = f_X$, reflecting the assumption that operating profits are non-increasing in t .

conventional sorting no longer holds: both the lowest- and the highest-cost active firms engage in FDI. Hence we can conclude that supermodularity of the profit function in t and c is necessary as well as sufficient for higher-cost firms to select into exports, and lower-cost firms to select into FDI, for all admissible fixed costs f .




Figure 3: Absent Supermodularity, the Conventional Sorting Fails for Some Fixed Costs

A striking feature of Proposition 1 is that it does not depend directly on fixed costs. While fixed costs affect the *level* of the tariff-jumping gain γ , they vanish when we compare across two firms using the finite difference operator Δ_c . Fixed costs are essential for a proximity-concentration trade-off, and hence they are necessary for the *existence* of selection effects. However, they do not determine their *direction*. So statements like “Only the more productive firms select into the higher fixed-cost activity” are often true, but always misleading: they are true given supermodularity, but otherwise not.¹³ What matters for the direction of selection effects is not a trade-off between fixed and variable costs, but whether there is a complementarity between variable costs of production and of trade. Putting this differently, for FDI to be the preferred mode of market access, a firm must be able to afford the additional fixed costs of FDI, but whether it can afford them or not depends on the cross-effect on profits of tariffs and production costs. When supermodularity prevails, a more efficient firm has relatively higher operating profits in

¹³The quoted statement is from Oxford graduate trade lecture notes in late 2009.

the FDI case, but when submodularity holds, the opposite is true.

Of course, all this assumes that fixed costs are truly fixed, both for a single firm as output varies, and for cross-section comparisons between firms. Matters are different if they depend on either t or c , as we shall see in Section 6. First, we turn to compare Proposition 1 with the result obtained by Helpman, Melitz, and Yeaple (2004). Our result is more general than theirs in that it places no restrictions on the functional form of the demand function: only mild restrictions on the maximized profit function $\pi(t, c)$ are needed. However, at first sight our result seems more special since it holds only for the case of a single monopoly firm, whereas Helpman, Melitz, and Yeaple (2004) proved their result in a general equilibrium model with monopolistic competition. In the next section we show that this apparent limitation of our result is illusory. With suitable reinterpretation, our result holds in a large class of monopolistically competitive models, including that of Helpman, Melitz, and Yeaple (2004).

4 Monopolistic Competition

To see how our result extends to models of monopolistic competition, we need to address the issue of market structure per se, to explore demand systems other than the CES, and to examine the specification of transport costs. We consider these issues in turn.

4.1 Exports versus FDI with CES Preferences

As already noted, Helpman, Melitz, and Yeaple (2004) were the first to consider how firms of different costs will select into different modes of serving foreign markets. They considered this issue in a model of heterogeneous firms in monopolistic competition with CES or Dixit-Stiglitz preferences, iceberg transport costs, and a Pareto distribution of firm productivities. In our notation the variable-profit function for a typical firm in such a model is:

$$\pi(t, c) = (\tau c)^{1-\sigma} B \tag{11}$$

where $\tau \equiv 1 + t \geq 1$ is an iceberg transport cost;¹⁴ σ is the elasticity of substitution in demand, which must be greater than one; and B is a catch-all term which summarizes the dependence of the demand for one firm's good on total expenditure and the prices of all other goods.¹⁵ Consider first the partial equilibrium or firm-level case where B is taken as given. In that case, it is clear that the profit function is supermodular in t and c :

$$\pi_{tc} = (\sigma - 1)^2 (\tau c)^{-\sigma} B > 0 \quad (12)$$

Hence, from Proposition 1, the ranking of firms by their mode of serving foreign markets established by Helpman, Melitz, and Yeaple (2004) follows immediately without any need to compare the levels of profits in different modes.

In full industry equilibrium, the demand term B is endogenous. It depends directly on the level of total expenditure E and on the overall price index P in the market in question, while P in turn depends on all the variables that affect the global equilibrium, including at a minimum the number of active firms serving this market from every country i , the distribution of firm costs $g(c)$, and the transport cost τ :

$$B = \tilde{B}(E, P) \quad P = \tilde{P}[\{n_i\}, g(c), \tau] \quad (13)$$

However, for the comparisons we wish to make, this endogeneity is not relevant. The price index and hence the demand term B would be affected by *changes* in transport costs which disturb the full equilibrium. But our concern is rather with characterizing the pattern of firm selection between different modes of serving foreign markets which face *differences* in transport costs. Since any pair of firms is infinitesimal relative to the mass of all firms, we can compare their choices while holding constant the actions of all other firms. Hence, partial equilibrium is the appropriate framework for the cross-section comparisons between different firms in the same equilibrium that we want to make.

This key point can be made differently by considering Figure 4, which is based on

¹⁴For continuity with previous sections we continue to write ex post profits as a function of t . This is not a restriction since $\partial\pi/\partial\tau = \pi_t$.

¹⁵In typical specifications, $B = (\sigma - 1)^{\sigma-1} A / \sigma^\sigma$, where A is the constant term in the demand function $x = Ap^{-\sigma}$. A in turn depends on nominal expenditure E and the aggregate price index P : $A = EP^\sigma$.




Figure 4: Inferring Selection Effects from Supermodularity

Helpman, Melitz, and Yeaple (2004). Their approach, now standard in the literature, is to compute the general equilibrium of the world economy and then to investigate what pattern of selection effects it exhibits. Thus they calculate not only the profit functions Π^F and Π^X , allowing for their dependence on expenditure and price indices in general equilibrium, but also their point of intersection, which is the threshold cost level at which a firm is indifferent between exports and FDI. By contrast, our approach is very different. We *assume* that an equilibrium exists, and that π is supermodular. We can then pick an arbitrary pair of firms, say those with the unit costs c_1 and c_2 in Figure 4. Rewriting the supermodularity condition $\Delta_c \pi(t, c) > \Delta_c \pi(0, c)$, and adding $-f_F + f_X$ to both sides gives a ranking of the two firms' total profits when they engage in FDI rather than exporting:

$$\begin{aligned}
 \pi(t, c_1) - \pi(t, c_2) &> \pi(0, c_1) - \pi(0, c_2) & (14) \\
 \Leftrightarrow \pi(0, c_2) - \pi(t, c_2) &> \pi(0, c_1) - \pi(t, c_1) \\
 \Leftrightarrow \Pi^F(c_2) - \Pi^X(c_2) &> \Pi^F(c_1) - \Pi^X(c_1)
 \end{aligned}$$

Repeating this comparison for every pair of firms allows us to infer the qualitative properties of the Π^F and Π^X loci without the need to calculate the full equilibrium. Our approach parallels that of Maskin and Roberts (2008), who show that all the central the-

orems of normative general equilibrium theory can be proved using elementary methods once the existence of equilibrium has been established. Our approach cannot confirm that an equilibrium exists. However, by dispensing with computing one explicitly, it applies without specific restrictions on the functional forms of preferences, technology, or the distribution of costs; it avoids the need to assume that countries are symmetric; and it extends easily to considering firm choices other than that between exports and FDI, as we shall see in Section 6.

4.2 General Preferences

While the result in the last sub-section has already been derived by Helpman, Melitz, and Yeaple (2004), the strength of our approach is that it allows us to sign selection effects into FDI for a much wider class of preferences than the CES. Write the demand function perceived by the firm in inverse form, $p = p(x)$, with no restrictions other than that consumers' willingness to pay is decreasing in price, $p' < 0$; and write the elasticity of demand as a function of sales: $\varepsilon(x) \equiv -\frac{\partial x}{\partial p} \frac{p}{x} = -\frac{p}{xp'}$. To determine which specifications of demand favor the conventional sorting, we introduce the term “superconvex” demand: we define a superconvex demand function as one for which $\ln p$ is convex or more so in $\ln x$.¹⁶ As we show in the Appendix, this is equivalent to the demand function being equally or more convex than a constant-elasticity CES demand function (for which ε equals σ), and to one whose elasticity of demand is increasing in output, so ε_x is non-negative. The case where demand is not superconvex, so ε is decreasing in x , we call subconvex, and, as noted by Krugman (1979), it may be considered the “normal” case: it implies plausibly that consumers are more responsive to price changes the greater their consumption; and it encompasses many of the most widely-used non-CES specifications of preferences, including the quadratic (to be considered further below) and the

¹⁶For a formal definition, and proofs of the statements that follow, see the Appendix, Section 8.2. The term “superconvexity” seems to be used, if at all, as a synonym for logconvexity, i.e., $\ln x$ convex in p . (See Kingman (1961).) For related discussions, see Neary (2009), Zhelobodko, Kokovin, Parenti, and Thisse (2011), and Dhingra and Morrow (2011). All these papers assume that preferences are additively separable, though this is not necessary for our approach, since we only consider the demand function from the firm's perspective.

translog.¹⁷ Strict superconvexity is less widely encountered; an example is where the inverse demand function has a constant elasticity relative to a displaced or “translated” level of consumption: $p = (x - \beta)^{-1/\sigma}$ with β strictly positive.¹⁸

The importance of superconvexity in this context is shown by the following result:

Proposition 2. *With general demands and iceberg transport costs, a sufficient condition for the profit function to be supermodular in t and c for all levels of output is that the demand function is weakly superconvex, i.e., the elasticity of demand is non-decreasing in output, $\varepsilon_x \geq 0$.*

The proof, given in the Appendix, follows by expressing the cross-partial derivative of the profit function in terms of the elasticity of demand and its responsiveness to output:

$$\pi_{tc} = \frac{(\varepsilon - 1)^2 + x\varepsilon_x}{\varepsilon - 1 - x\varepsilon_x}x \quad (15)$$

The denominator $\varepsilon - 1 - x\varepsilon_x$ must be positive from the second-order condition. Hence the numerator shows that, whenever ε_x is strictly negative, submodularity may hold for sufficiently high x . We can be sure that supermodularity holds for all output levels only in the CES and strictly superconvex cases.¹⁹

Intuitively, the result follows from another implication of superconvexity. A positive value of ε_x means that larger firms face a higher elasticity of demand. Since output is decreasing in c in this model ($x_c < 0$), this implies that, if and only if ε_x is positive, more productive firms face more elastic demand. Hence, they also have lower mark-ups, as measured by the Lerner Index, $L \equiv \frac{p-\tau c}{p}$, since $L = \frac{1}{\varepsilon}$.²⁰ This implies that a

¹⁷Feenstra (2003) shows how the translog can be adapted to allow for a variable number of varieties consumed, and so used in models of monopolistic competition with free entry. He also shows that it implies an elasticity of demand which is always increasing in price, and so, from Lemma 7 in the Appendix, is subconvex. (See Feenstra (2003), p. 85).

¹⁸We are grateful to Rob Feenstra for suggesting this example.

¹⁹In the CES case, when ε_x is zero, (15) reduces to $\pi_{tc} = (\sigma - 1)x$ and is always positive. This is equivalent to equation (12), using the fact that output with CES preferences equals: $x = \left(\frac{\sigma}{\sigma-1}\tau c\right)^{-\sigma}$.

²⁰It can be checked that the Lerner index falls as costs fall if and only if ε_x is positive: $\frac{dL}{dc} = -\frac{\varepsilon_x}{\varepsilon^2} \frac{dx}{dc}$. This result does not hold for the *absolute* mark-up, $\mu \equiv p - \tau c = \frac{p}{\varepsilon}$, which can be either increasing or decreasing in production cost. For example, it is increasing in c in the CES case: $\mu = \frac{\tau c}{\sigma-1}$, and so $\frac{d\mu}{dc} = \frac{\tau}{\sigma-1} > 0$; but decreasing in c in the quadratic case to be discussed further below: $\mu = \frac{A-\tau c}{2}$, and so $\frac{d\mu}{dc} = -\frac{\tau}{2} < 0$.

more productive firm will have an incentive to expand output more in order to maximize profits. As a result, the Matthew Effect is stronger when ε_x is positive, sufficiently so that supermodularity is guaranteed. By contrast, when ε_x is negative, the Matthew Effect is weaker and so more productive firms may not benefit as much from avoiding the tariff by engaging in FDI.

Proposition 2 is important for highlighting which classes of demand function are consistent with super- or submodularity, but it is only a sufficient condition. To determine whether a particular demand function exhibits supermodularity, we can use the necessary and sufficient condition given by the following:

Proposition 3. *With general demands and iceberg transport costs, a necessary and sufficient condition for the profit function to be supermodular in t and c is that the sum of the elasticity and convexity of demand is greater than three.*

The proof (given in the Appendix) proceeds by showing that equation (15) can be reexpressed as follows:

$$\pi_{tc} = \frac{\varepsilon + \rho - 3}{2 - \rho}, \quad \rho \equiv -\frac{xp''}{p'} \quad (16)$$

where ρ is our measure of convexity of demand. So, submodularity is more likely when demand is less elastic and more concave. In particular, it may arise for any linear or concave demand system, and even for demands that are “not too” convex.

To illustrate how these results can be applied in practice, we consider two subconvex demand systems, one of which always implies that profits are supermodular and the other which implies submodularity for high values of output. The first is a version of the translated CES case already mentioned:

Lemma 2. *The demand function $p = (x - \beta)^{-1/\sigma}$, with β strictly negative, is always subconvex, but the implied profit function is supermodular at all levels of output provided $\sigma \geq 1.25$.*

In this case, higher sales are associated with a lower demand elasticity and thus a higher markup, implying that more productive firms do not exhibit such a large difference in

output. Nevertheless, the elasticity of demand never falls sufficiently low to allow submodularity to emerge.²¹

Our second example is the case of quadratic preferences, which have been studied in the context of heterogeneous firms by Melitz and Ottaviano (2008). Their model has been extended to the choice between exports and FDI by Nefussi (2006), but only by solving for the full general equilibrium. Using our approach it is easy to establish its properties and to show that its predictions for firm selection into FDI are ambiguous. Writing the inverse demand function as $p = A - bx$, we can express the elasticity of demand as a function of output:²²

$$\varepsilon(x) = \frac{A - bx}{bx} \quad (17)$$

This is monotonically decreasing in x , so from Proposition 2 we know that for high values of output the profit function may be submodular. However, we need to check that this will happen for values of x that are admissible in equilibrium. To confirm this we specialize Proposition 2 to the quadratic case:

Lemma 3. *With quadratic preferences and iceberg transport costs, the profit function π is supermodular when defined on the interval $c \in [\frac{A}{2\tau}, \infty]$ and submodular when defined on the interval $c \in [0, \frac{A}{2\tau}]$.*

(The proof is in the Appendix.) Unlike the CES case, the profit function is therefore *submodular* for low-cost exporters, although it continues to be *supermodular* for high-cost ones. Hence, provided both exporting and FDI are profitable in the relevant range, we can expect a threefold selection effect in this model: the very highest-cost firms will select into exporting, but so will the very lowest-cost ones, while intermediate-cost firms will select into FDI.²³ Figure 5 illustrates this configuration.

²¹As shown in the Appendix, Section 8.5, $\varepsilon = \frac{x-\beta}{x}\sigma$ and $\rho = \frac{x}{x-\beta}\frac{\sigma+1}{\sigma}$. Though the demand function is subconvex for $\beta < 0$, it is always strictly convex. It follows that $\varepsilon + \rho - 3 = \varepsilon + \frac{\sigma+1}{\varepsilon} - 3$, which can be negative only for very low σ .

²²As always in monopolistic competition, the demand parameters A and b are taken as given by firms, but are endogenous in general equilibrium. For example, in the Melitz-Ottaviano framework, $A \equiv \frac{\gamma\alpha+\eta N\bar{p}}{\gamma+\eta N}$ and $b \equiv \frac{\gamma}{L}$, where α , γ and η are demand parameters, L is market size, N is the mass of firms, and \bar{p} is the aggregate price index.

²³This case holds provided a number of boundary conditions are met: (i) exporting must be profitable, $\Pi^X \equiv \pi(t, c) - f_X > 0$, which requires: $c < \frac{1}{\tau}(A - 2\sqrt{bf_X})$; (ii) FDI must be profitable, $\Pi^F \equiv$




Figure 5: Selection Effects with Quadratic Preferences and Iceberg Transport Costs

As already noted, the corollary given here generalizes the result of Nefussi (2006), dispensing with the assumptions of symmetric countries and a Pareto distribution of firm productivities which he makes. Our result also extends easily to explain the pattern of firm selection into exporting, export-platform FDI, and multi-market FDI when there is more than one foreign market. This is particularly convenient because, as Behrens, Mion, and Ottaviano (2010) show, it does not seem to be possible to compare two different FDI equilibria analytically when preferences are quadratic. The problem arises from the fact that all variables in any given equilibrium can be written as functions of the cost cutoff (the threshold level of marginal cost above which a firm finds it unprofitable to produce).²⁴ However, comparing two different cutoffs is extremely difficult. Our approach makes it unnecessary to do so: we *assume* that an equilibrium exists in which firms select into different modes of serving the market, and can then invoke our result on supermodularity

$\pi(0, c) - f_F > 0$, which requires: $c < A - 2\sqrt{bf_F}$; and (iii) some selection must take place, i.e., the quadratic equation in c defined by $\Pi^X = \Pi^F$ must have two real roots, which requires: $(\tau - 1)A^2 > 4(\tau + 1)b(f_F - f_X)$. Note that we allow for a non-zero fixed cost of exporting, unlike Melitz and Ottaviano (2008). To solve their model in full, they have to assume that exports do not incur any fixed costs, in which case the demand parameter A equals the marginal cost of the threshold firm in equilibrium. Our approach can accommodate fixed costs of exporting, so this property does not necessarily hold here.

²⁴Though this is only possible if there is no fixed cost of exporting, an assumption which our approach does not require, as discussed in the previous footnote.

to justify which mode is relatively more profitable for any pair of firms, and, by extension, for all firms.

4.3 General Transport Costs

The result in the previous sub-section that the largest firms select into exporting for a wide class of demand functions is not necessarily paradoxical. It may simply be viewed as yet another example of large firms’ “supermodular superiority”.²⁵ To the extent that the most efficient firms are more productive in all the activities in which they engage, then it is reasonable to assume that they also incur the lowest per unit transport costs. Perhaps they are able to avail of economies of scale in transportation, or to negotiate better terms with transport contractors. From that perspective, the assumption of iceberg transport costs can be seen as a convenient reduced-form way of modeling this superiority of more efficient firms. On the other hand, the suspicion remains that this result is an artifact of iceberg transport costs. It is stretching credulity to assume that the most efficient firms produce the cheapest icebergs, and, in particular, that highly efficient firms, with production costs close to zero, also incur negligible transport costs irrespective of distance. But this is what is implied by the iceberg assumption: to sell q units it is necessary to produce and ship τq units, so the technology of transportation is identical to that of production: $(p - \tau c)x = px - c(\tau x)$.

To see how alternative specifications of transport costs affect the outcome, consider a general specification of the ex post profit function as the outcome of choosing output x to maximize $\tilde{\pi}(x; \tau, c)$, the firm’s operating profits as a function of the exogenous variables τ and c and an arbitrary level of output:

$$\pi(t, c) \equiv \underset{x}{Max} \tilde{\pi}(x; \tau, c) \tag{18}$$

We can now express the desired cross-partial derivative of π in terms of second derivatives

²⁵We are grateful to Adrian Wood for suggesting this line of reasoning.

of $\tilde{\pi}$:²⁶

$$\pi_{tc} = \tilde{\pi}_{\tau c} + \tilde{\pi}_{\tau x} \frac{dx}{dc} = \tilde{\pi}_{\tau c} - \tilde{\pi}_{\tau x} (\tilde{\pi}_{xx})^{-1} \tilde{\pi}_{xc} \quad (19)$$

This shows that supermodularity of the profit function in t and c depends on the balance between two effects: a direct effect given by $\tilde{\pi}_{\tau c}$, which is the effect of a difference in production costs on the profit disadvantage of higher transport costs at a *given* level of output; and an indirect effect given by the second term on the right-hand side. The expression $\tilde{\pi}_{xx}$ is negative from the firm's second-order condition, so the sign of the indirect effect depends on the product $\tilde{\pi}_{\tau x} \tilde{\pi}_{xc}$. This is presumptively positive; for example it must be so in the case of constant production costs and iceberg transport costs, when $\tilde{\pi}_{\tau x} = -c$ and $\tilde{\pi}_{xc} = -\tau$. This is the Matthew Effect from Section 2: it arises because a higher-cost firm is less vulnerable to a rise in transport costs since it has presumptively lower sales: both $\tilde{\pi}_{\tau x}$ and $\frac{dx}{dc}$ are negative, so their product is positive. By contrast, the direct effect is less robust. In the case of iceberg transport costs it simply equals $\tilde{\pi}_{\tau c} = -x$ and is clearly the source of the potential for submodularity identified in the previous sub-section. It reflects the fact that a higher-cost firm loses more from a rise in transport costs ($\tilde{\pi}_{\tau}$ is more negative) since its cost of shipping one unit of exports is $(\tau - 1)c$.

It is immediate that the direct effect vanishes if transport costs and production costs are separable in the profit function $\tilde{\pi}$. This corresponds to the case where exports do not melt in transit, but trade costs are levied instead on the value of sales:

$$\tilde{\pi}(x; \tau, c) = R(x, \tau) - cx \quad (20)$$

Here net sales revenue accruing to the firm, R , depends in a very general way on the transport cost parameter. However, there is no interaction between transport costs and production costs. As a result, there is no direct effect in the supermodularity expression given by (19): total transport costs and hence $\tilde{\pi}_{\tau}$ do not depend directly on c , implying that the direct effect $\tilde{\pi}_{tc}$ is zero. By contrast, the indirect effect is positive as before.

²⁶To derive this we use the envelope theorem to set $\pi_t = \tilde{\pi}_{\tau}$, and totally differentiate the first-order condition $\tilde{\pi}_x = 0$ to obtain $\frac{dx}{dc} = -(\tilde{\pi}_{xx})^{-1} \tilde{\pi}_{xc}$.

Hence, profits are supermodular in t and c for all levels of output and all specifications of demand when transport costs and production costs are separable in this way.

Specific transport costs, already considered in Example 1 in Section 2 above, provide one example of (20). Another is where transport costs are ad valorem or proportional to price, so net sales revenue becomes: $R(x, \tau) = \frac{xp(x)}{\tau}$. Relative to the case of iceberg transport costs, the firm's first-order condition is unchanged, but profits are deflated by τ : $\tilde{\pi}(x; \tau, c) = \left[\frac{p(x)}{\tau} - c \right] x$. Similar derivations to those already given shows that equation (19) now becomes: $\pi_{tc} = -c(2p' + xp'')^{-1} > 0$. Thus the full effect is unambiguously positive for *all* demand systems, and so the profit function is always supermodular. Figure 6 illustrates the case of quadratic preferences and proportional transport costs. Clearly, the conventional sorting is now restored, and the model predicts that the most efficient firms will always engage in FDI rather than exporting.




Figure 6: Selection Effects with Quadratic Preferences and Ad Valorem Transport Costs

5 Selection Effects in Oligopoly

The previous section, like almost all the recent literature on trade with heterogeneous firms, assumed that markets are monopolistically competitive. Rare exceptions to this generalization include Porter (2011), who shows that the more efficient firm in a duopoly

is more likely to engage in FDI than exporting, and Leahy and Montagna (2009) who show a similar result for outsourcing. It is desirable to establish whether similar results hold more generally when firms are large enough to exert market power over their rivals, so markets are oligopolistic. As already noted, this is of interest both as a check on the robustness of the results and also because, to the extent that more successful firms are likely to engage in a wider range of activities, the assumption that they remain atomistic relative to their smaller competitors becomes harder to sustain.

If individual firms are no longer of measure zero then the arguments used in Section 4.1 no longer hold. If we wish to compare a firm's profits under exporting and FDI, we can no longer assume that the industry equilibrium is unaffected by its choice. However, our earlier result still holds when we take behavior by rival firms as given. To illustrate with a simple example, consider the case where there are two rival U.S. firms, labeled "1" and "2", both of which consider the choice between exporting to the EU and locating a foreign affiliate there. The payoffs to firm 1, conditional on different choices of firm 2, are given in Table 1. Thus, the first entry in the first row, $\pi(t, c, X) - f_X$ gives the operating profits which it will earn if it exports to the foreign market, conditional on the rival firm 2 also exporting. We would expect this to be always less than the second entry, $\pi(t, c, F) - f_X$, which is conditional on firm 2 engaging in FDI: better market access by the rival presumably reduces firm 1's profits, *ceteris paribus*. However, what matters for firm 1's choice is the comparison between different entries in the same column, and it is clear that, conditional on a given mode of market access by firm 2, firm 1's choice will reflect exactly the same considerations as in previous sections. Hence, provided supermodularity holds in each column, and in the columns of the corresponding table for firm 2, our earlier result goes through: when that is the case, more efficient firms will select into FDI and less efficient ones into exporting.

While the central result derived earlier still holds, it has to be applied with care. One issue is that boundary cases have to be considered in detail. Depending on the configuration of the two firms' costs, in the Nash equilibrium only one of them may serve the market at all, or do so via FDI. There may be no equilibria in pure strategies, in

Choice of Firm 2:	Export	FDI
Export:	$\pi(t, c, X) - f_X$	$\pi(t, c, F) - f_X$
FDI:	$\pi(0, c, X) - f_F$	$\pi(0, c, F) - f_F$

Table 1: Payoffs to Firm 1 Given Choices of Firm 2

which case mixed strategy equilibria have to be considered. Finally, the necessity part of Proposition 1 does not survive. This is because, even when we allow for all values of fixed costs as in the proof of Proposition 1, supermodularity of the profit function conditional on rivals' responses is necessary for the conventional sorting only at those points which are relevant to a particular Nash equilibrium. Thus it is conceivable that supermodularity might not hold over a range of the profit function; but if that range was never relevant for any value of fixed costs, then the conventional sorting would still apply.

6 Other Applications

6.1 Alternative Firm Choices

So far, our focus has been on the choice between exports and FDI. However, the same approach applies to a wide range of other firm choices. For example, we show in Mrázová and Neary (2010) that similar considerations can be used to predict whether firms will select into export-platform or multi-plant FDI, when they are faced with the problem of how to serve a number of foreign markets. We also show below that our approach extends to the choice between in-house production and outsourcing, as in Antràs and Helpman (2004), and to the choice between more and less skill-intensive technology, as in Bustos (2011). In these and many other cases, results analogous to those derived above apply: supermodularity between the firm's own cost parameter and a parameter representing the marginal cost of the mode of accessing a market is necessary and sufficient for the standard selection effect, whereby more productive firms select into the access mode with lower marginal cost.

One key exception to our framework is the “zero-one” choice between serving a market or not, and a central example of this case is the decision on whether to engage in exporting or not, as originally considered by Melitz (2003). Selection effects in that case depend only on the *level* of total profits attainable by exporting, $\Pi^X \equiv \pi(t, c) - f_X$, rather than on the *difference* between the profits attainable via exports, Π^X , and those attainable via FDI, Π^F . Selection effects between firms with different productivities thus depend in the Melitz case on the *first* derivative of the variable profit function rather than on the more subtle *second cross*-derivative which we have emphasized: π decreasing in c ensures the conventional sorting whereby there is a single threshold level of costs such that all firms with lower costs will export and all those with higher costs will serve the home market only. It is perhaps appropriate that, unlike many of its applications and extensions, the model of Melitz makes a much more robust prediction about selection effects, one which does not depend on the properties of the second derivative of the profit function.

In the remainder of this section, we present some applications of our approach. We first look at the issues of both choice of location and choice of organizational form which arise when firms can vertically disintegrate. We then turn to show how our approach extends to the case where fixed costs differ between locations and between firms. Finally, we consider how selection effects can also be inferred in models where fixed costs are endogenous, determined by prior investments in variables such as technology, research and development (R&D), or marketing.

6.2 Vertical Disintegration and Choice of Organizational Form

Our discussion of FDI in previous sections concentrated on the horizontal kind, where the firm is considering how to serve a foreign market, and FDI involves effectively reproducing abroad the production facilities which are already located in the home country. This archetypal problem is also one in which differences between the two countries are not central: in particular, we assumed for convenience that the marginal cost of production was the same whether the firm engaged in exporting or in FDI. A different problem arises in the case of a firm whose goal is to serve its home market, but which faces two distinct

choices about its organizational form. On the one hand, it has the option of producing either at home or in a lower-cost location abroad. On the other hand, it can choose either to produce in-house or to outsource: the choice of whether or not to vertically integrate arises irrespective of where production is located. The classic treatment of this issue in a model with heterogeneous firms is by Antràs and Helpman (2004), and we draw on their work in what follows. However, since our specification is a reduced-form one, it is also consistent with many other ways of modeling the choice of organizational form.

Ignoring fixed costs for the present, the choice of production location and organizational form will depend on the total operating profits that are realized in each case. We write this as:

$$\pi(w, \psi, c) \equiv \underset{x}{Max} (1 - \psi) [p(x) - wc] x \quad (21)$$

Unlike in previous sections, we abstract from transport costs. In other respects the model is more complicated. First, the wage which the firm must pay differs between locations. If production takes place in the home country, which we will henceforth refer to as “North”, the wage equals w_N , while if it takes place in “South” it is lower and equal to w_S , with $w_S < w_N$. Second, the firm owner must use the services of the supplier of an intermediate input, the quality of which, though observable to both parties, is not contractible.²⁷ This leads to a profit loss due to incomplete contracting between the firm owner and the intermediate-input supplier, represented by the parameter ψ . Structural microfoundations for this parameter are provided in Antràs and Helpman (2004).²⁸ Here we need only assume that it may differ both between location and organization form; in particular, we assume that, irrespective of location, it is lower when the firm vertically integrates than when it outsources and must contract with an outside supplier: $\psi_j^V < \psi_j^O$, $j = N, S$.²⁹

²⁷Either the costs of writing a comprehensive contract are infinite, or the outcomes cannot be observed by a third-party arbitrator. In either case a complete contract cannot be enforced.

²⁸They assume that the input supplier’s outside option is zero, so it must be paid a fixed amount to persuade it to participate. The headquarters maximizes $(1 - \psi) \pi - T - f$ where T is the fixed amount that it must pay the supplier. This is equivalent to maximizing $(1 - \psi) \pi - f$.

²⁹Our reduced-form specification covers both the case where internalization eliminates all costs of incomplete contracting, as in Williamson (1975) and Grossman and Helpman (2002), and the case where even vertical integration in the North incurs some cost, as in Grossman and Hart (1986), Hart and Moore (1990), and Antràs and Helpman (2004). Our approach can easily be extended to allow the efficiency

Considering only variable costs, the firm owner has an incentive to locate production in the South, and to produce in-house rather than outsource. Offsetting these differences in variable costs, the fixed costs of locating in the South are higher than in the North, irrespective of the choice of organizational form: $f_S^i > f_N^i$, $i = O, V$; and the fixed costs of vertical integration or in-house production are higher than those of outsourcing, irrespective of the choice of location: $f_j^V > f_j^O$, $j = N, S$. This configuration of fixed and variable costs ensures that there is a trade-off between different modes of organization, and opens up the possibility that some selection by firms will take place. However, just as in previous sections, it does not predict the direction of selection effects, except in the special case of CES preferences considered by Antràs and Helpman (2004).

Firm's Decision	Profits	$\varepsilon_x \geq 0$	$\varepsilon_x < 0$
HFDI with iceberg transport costs:	$\pi(t, c) = (p - \tau c)x$	Yes	No
HFDI with proportional transport costs:	$\pi(t, c) = (p/\tau - c)x$	Yes	Yes
Produce in North or South:	$\pi(w, c) = (p - wc)x$	Yes	No
Produce in-house or outsource:	$\pi(\psi, c) = (1 - \psi)(p - wc)x$	Yes	Yes

Table 2: Is the Profit Function Supermodular at all Levels of Output?

To see this, we consider in turn the two choices which the firm must make. Consider first the simplest form of the choice of location, where we assume that there is no efficiency cost of incomplete contracting, so ψ in (21) equals zero. This corresponds to the choice between outsourcing to a Northern contractor and offshoring to a Southern contractor in Antràs and Helpman (2004). Comparing the first and third row of Table 2, it is clear that the profit function in this case has exactly the same form as the profit function with iceberg transport costs in the horizontal FDI case, with the wage w playing the role of the iceberg cost parameter τ . Hence the results of previous sections apply immediately. From Proposition 1, if selection takes place, then more efficient firms will offshore and less efficient ones will outsource at home if and only if the profit function is supermodular in w and c ; while from Proposition 2 we can be sure that the profit function is supermodular

cost of incomplete contracting to depend on either firm productivity or wages or both: just replace ψ in (21) by $\Psi(\psi, c, w)$, where ψ now represents a structural parameter that determines the cost of incomplete contracting for given values of c and w .

for all output levels if and only if the elasticity of demand is constant or increasing in sales. Thus, without the need for any further analysis, we can conclude that more efficient firms select into outsourcing at home when preferences are CES, but not necessarily otherwise. For example, the sorting predicted by Antràs and Helpman (2004) need not hold under quadratic or translog preferences, even with all of their other assumptions retained: very efficient firms are likely to choose to outsource at home rather than to offshore, since they are less affected by the higher wage in the North. On the other hand, with CES preferences, their results go through irrespective of how firm productivities are distributed or whether countries are symmetric or not, provided only that an equilibrium exists.³⁰

Consider next the simplest type of choice of organizational form, where we hold location constant and hence the wage is given. From the fourth row of Table 2, selection effects now depend on whether the profit function is supermodular in ψ and c . It turns out that this is always the case. The direct effect of a difference in costs between firms is straightforward: the higher-cost firm has lower profits, and as a result is less affected by a higher profit leakage due to incomplete contracts. In fact, the direct effect is simply $\tilde{\pi}_{\psi c} = x$. As for the indirect effect, it equals zero: although the higher-cost firm has lower sales, it is neither more nor less vulnerable to a rise in the profit leakage parameter ψ , since this is analogous to a uniform profits tax on all firms: $\tilde{\pi}_{\psi x} = -\tilde{\pi}_x = 0$. Thus, adapting equation (19) to this case, we have:

$$\pi_{\psi c} = \tilde{\pi}_{\psi c} - \tilde{\pi}_{\psi x} (\tilde{\pi}_{xx})^{-1} \tilde{\pi}_{xc} = x > 0 \quad (22)$$

This implies that the selection effect for outsourcing is very robust: more efficient firms select into the lower- ψ organizational form (i.e., producing in-house rather than outsourcing) for *all* preferences and productivity distributions.

Finally, selection effects into vertical FDI are ambiguous in general. Relative to pro-

³⁰Note a further generalization which our approach makes possible. Antràs and Helpman (2004) assume that fixed costs are always incurred in the North, so $\Pi = \pi - w_N f$. Suppose instead that fixed costs are incurred in the country where production is located. In that case we can redefine variable profits as $\pi(w, c) = (p - wc)x - wf$, $w = w_N, w_S$, and we can see that this makes no difference to whether π is supermodular or not in w and c . Hence the assumption made by Antràs and Helpman is not needed for the results. (The only complication is that, for the necessity part of the proof of Proposition 1, we need to assume that there is some component of fixed costs which is independent of w .)

ducing as an integrated firm at home, vertical FDI faces conflicting incentives. On the one hand, it incurs a higher cost of incomplete contracting, $\psi_S^V > \psi_N^V$, since in the event of the relationship breaking down, the headquarters cannot expect to retain as large a share of profits when production is in the South. On the other hand, it incurs a lower wage rate, $w_S < w_N$. Thus the profit function is unlikely to be unambiguously supermodular or submodular in c and the cost vector $[\psi \ w]$, and the predicted selection effects are likely to be highly sensitive to the specification of the model.

6.3 Heterogeneous Fixed Costs

Up to this point we have followed most of the literature on heterogeneous firms in assuming that fixed costs are the same for all firms and in all foreign locations. This is clearly unrealistic, and we need to examine whether our approach can be extended to the case where fixed costs differ between firms or locations.

The previous analysis is unaffected if fixed costs vary with trade costs t only, so f becomes $f(t)$. For example, Kleinert and Toubal (2010) allow the fixed costs of a foreign plant to increase with its distance from the parent country, and show that this change in assumptions rationalizes a gravity equation for FDI, while Kleinert and Toubal (2006) show that it also avoids the counter-factual prediction that falling trade costs lower FDI. These are important insights, but the model's predictions about selection effects are unchanged. The reason is simple: although the fixed cost varies with trade costs, the finite difference operator applied to the gain from FDI relative to exporting γ eliminates the fixed cost since $\Delta_c f_F(t) = 0$. While differences in fixed costs between locations clearly affect locational choice, they do so in the same way for all firms.

Matters are more complicated if fixed costs vary with both production costs c and trade costs t . Technically, our approach can still be applied, but some care is needed. We now need to include any firm-specific fixed costs in the definition of operating profits.³¹ Thus, let $\tilde{\pi}(t, c)$ denote operating profits net of firm-specific fixed costs, and define full

³¹If the necessity part of Proposition 1 is to hold, we must also assume that there is a component of fixed costs which is common to all firms, as before.

operating profits $\pi(t, c)$ as follows:

$$\pi(t, c) \equiv \tilde{\pi}(t, c) - I(t)f_F(c) \quad (23)$$

Here, $I(t)$ is an indicator function, which equals zero when t is positive (the exporting case), and equals one when t is zero (the FDI case). Thus the fixed cost of a foreign plant must be subtracted to get full operating profits, and this varies with the firm's productivity. Now, there is an additional reason why supermodularity may not hold, depending on how fixed costs vary with productivity. Applying the finite difference operator to the full operating profits function (23) gives:

$$\Delta_c \pi(t, c) - \Delta_c \pi(0, c) = [\Delta_c \tilde{\pi}(t, c) - \Delta_c \tilde{\pi}(0, c)] + \Delta_c f_F(c) \quad (24)$$

The first term in parentheses on the right-hand side is the same as in previous sections. The second term is new, and shows that supermodularity is more likely to hold if fixed costs are higher for less efficient firms.

Two examples from the literature illustrate how this effect can work in different directions. The first is from Behrens, Mion, and Ottaviano (2010), who assume that a firm's fixed costs are proportional to its variable costs, $f_F(c) = cf$, so more efficient firms incur lower fixed costs of establishing a foreign plant. In this case, the final term in (24) becomes: $\Delta_c f_F(c) = (c_1 - c_2)f$ which is strictly positive for $c_1 > c_2$. Hence, supermodularity of π and so the conventional sorting pattern are reinforced in this case.

A second example comes from Oldenski (2010), who develops a model of task-based trade in services. Because they use knowledge-intensive tasks disproportionately, higher-productivity firms in service sectors are more vulnerable to contract risk when located abroad. As a result, their fixed costs of FDI are decreasing in c : $f_F(c), f'_F < 0$. As a result, the final term in (24) becomes $\Delta_c f_F(c) = f_F(c_1) - f_F(c_2)$, which is strictly negative for $c_1 > c_2$, so π may be submodular. In this case the conventional sorting may be reversed, as higher-productivity firms may find it more profitable to locate at home, and Oldenski presents evidence for this pattern in a number of U.S. service sectors.

6.4 Endogenous Fixed Costs

The previous sub-section considered fixed costs that differ between firms but in an exogenous way. By contrast, there are many ways in which a firm can influence the level of its fixed costs as well as its variable costs in each market. It is desirable to explore whether our approach extends to this case, where the firm faces more complex trade-offs.

In fact, the earlier derivations go through with relatively little modification. All that is needed is to redefine the maximized profit function as the outcome of the firm's choice of *both* its output and its level of investment. To fix ideas, consider the case of investment in cost-reducing R&D. (Other forms of investment, such as in marketing or product innovation, can be considered with relatively minor modifications.) Let k denote the level of investment which the firm undertakes. For simplicity we concentrate on examples where k is a continuous variable. However, it is clear that our approach does not require this, and that supermodularity also applies to cases of a discrete choice between a finite number of techniques, as in Bustos (2011).

Turning to specifics, we assume that investment in R&D incurs an endogenous fixed cost $F(k)$ but reduces average production costs, now denoted $C(c, k)$. Here c is, just as in earlier sections, a parameter representing the firm's exogenous level of costs (the inverse of its productivity), while k is chosen endogenously. $C(c, k)$ is increasing in c and decreasing in k , while fixed costs $F(k)$ are increasing in k . The maximum profits which the firm can earn in a market, conditional on t and c , are:

$$\pi(t, c) \equiv \max_{x, k} [\{p(x) - C(c, k) - t\}x - F(k)] \quad (25)$$

Unlike earlier examples, the firm now has two choice variables rather than one. It can be checked that π is supermodular in t and c for many commonly used specifications of the cost functions $F(k)$ and $C(c, k)$, so all our results apply in those cases too. However, there are also economically interesting examples where supermodularity is violated, and so the selection pattern of firms into different modes of serving foreign markets given by Proposition 1 is reversed.

To check whether the profit function (25) exhibits supermodularity in t and c , we proceed as in Example 1. The envelope theorem still applies, so the derivative of maximum profits with respect to the tariff equals minus the level of output: $\pi_t = -x(t, c)$. Hence it follows as before that: $\pi_{tc} = -x_c$. So, to check for supermodularity, we need only establish the sign of the derivative of output with respect to the cost parameter c . The first-order conditions for output x and investment k are:

$$p - C - t + xp' = 0 \quad (26)$$

$$-xC_k - F' = 0 \quad (27)$$

Totally differentiate these and write the results as a matrix equation:

$$\begin{bmatrix} 2p' + xp'' & -C_k \\ -C_k & -(xC_{kk} + F'') \end{bmatrix} \begin{bmatrix} dx \\ dk \end{bmatrix} = \begin{bmatrix} C_c dc + dt \\ xC_{kc} dc \end{bmatrix} \quad (28)$$

From the firm's second-order conditions, the diagonal terms in the left-hand coefficient matrix must be negative, and the determinant of the matrix, which we denote by D , must be positive.³² Solving for the effect of the cost parameter on output and substituting into π_{tc} gives:

$$\pi_{tc} = -x_c = D_+^{-1} \left[\underbrace{C_c (xC_{kk} + F'')}_{+} - xC_k C_{kc} \right] \quad (29)$$

The first term inside the brackets is positive from the second-order conditions, which works in favor of supermodularity of π . The second could work either way. In particular, the term could be negative, and so supermodularity might not prevail, if C_{kc} is negative, so a lower-productivity firm benefits more from investment, in the sense that its costs fall by more; or, equivalently, if C_{ck} is negative, so investment lowers the cost disadvantage of a lower-productivity firm.

In some of the most widely-used models of R&D, the ambiguity of the sign of (29) can be resolved. We consider three examples in turn.

³² $D = -(2p' + xp'')(xC_{kk} + F'') - C_k^2$

6.4.1 Linear-Quadratic Costs

Following d'Aspremont and Jacquemin (1988), a widely-used specification of R&D costs, also applied to the study of FDI by Petit and Sanna-Randaccio (2000), assumes that the marginal cost function is linear while the investment cost function is quadratic in k .³³

$$C(c, k) = c_0 - c_1^{-1}k \quad F(k) = \frac{1}{2}\gamma k^2 \quad (30)$$

Firms may differ in either the c_0 or c_1 parameters, but it is clear that in either case output must be decreasing in c : C_{ck} is zero if firms differ in c_0 and positive if they differ in c_1 .³⁴ Hence, the right-hand side of (29) is positive and supermodularity is assured for this specification of R&D costs.

6.4.2 Exponential Costs of R&D

An implausible feature of the d'Aspremont-Jacquemin specification is that the returns to investing in R&D are constant.³⁵ A more attractive and only slightly less tractable alternative due to Spence (1984) is also widely used:³⁶

$$C(c, k) = c_0 + c_1 e^{-\theta k} \quad F(k) = k \quad (31)$$

In this case investment lowers marginal production costs ($C_k = -\theta c_1 e^{-\theta k} < 0$) but at a diminishing rate ($C_{kk} = \theta^2 c_1 e^{-\theta k} > 0$), as illustrated in Figure 7(a) (drawn for $c_0 = \theta = 1$); while the direct cost of investment increases linearly in k ($F'' = 0$). Once again, firms may differ in either the c_0 or c_1 parameters, and supermodularity is assured if they differ in c_0 . However, matters are different if firms differ in c_1 (so we set $c_1 = c$ from now on). In this case, a lower-productivity firm benefits more from investment: $C_{ck} = -\theta e^{-\theta k} < 0$,

³³d'Aspremont and Jacquemin and Petit and Sanna-Randaccio also allowed for spillovers between firms.

³⁴When $c = c_1$, $C_{ck} = c^{-2} > 0$.

³⁵The linearity of C in k also suggests that the cost of production can become negative, though second-order conditions ensure that this never happens in equilibrium.

³⁶These specifications of $C(c, k)$ and $F(k)$ come from Section 5 and from equation (2.3) on page 104 of Spence (1984), respectively.

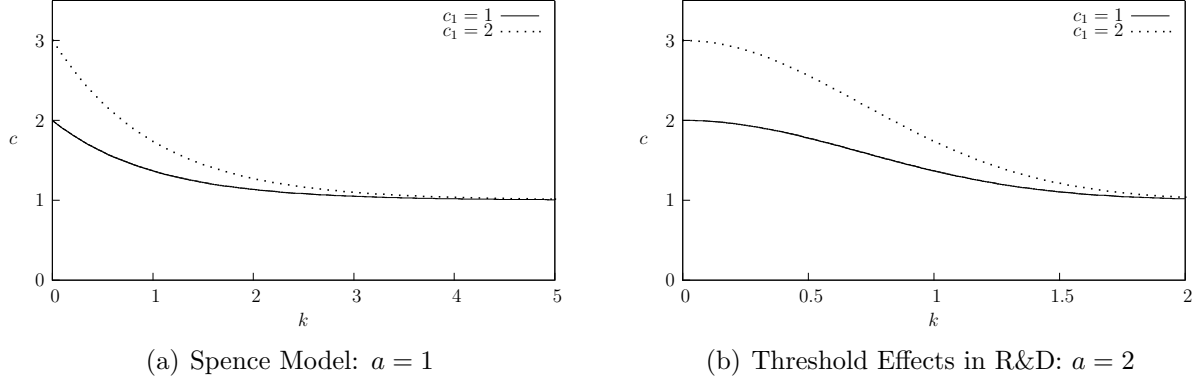


Figure 7: Marginal Cost of Production as a Function of Investment

and this effect is sufficiently strong that it exactly offsets the diminishing returns to investment.³⁷ As a result, equation (29) is zero under the Spence assumptions: other things equal, two firms with different cost parameters produce the same output.³⁸ Hence it follows that $\pi(t, c)$ is *modular*, i.e., both supermodular and submodular: the expression in Definition 1 holds with equality.

To see the implications of this for the decision to engage in FDI, consider how two firms of different productivities will assess the relative advantages of exporting and FDI. For any given mode of accessing a market, both firms will produce the same output, the less productive firm compensating for its higher *ex ante* cost by investing more.³⁹ Hence their operating profits are also the same. It follows that both firms face exactly the same incentive to export or engage in FDI. We cannot say in general which mode of market access will be adopted, but we can be sure that both firms will always make the same choice. More generally, for any number of firms that differ in c_1 , all firms will adopt the same mode of serving the foreign market, so no differences in firm selection will be observed.

³⁷Formally, the semi-elasticities of both C_c and C_k with respect to k , C_{ck}/C_c and C_{kk}/C_k , are equal to $-\theta$.

³⁸Evaluating (29) gives: $-C_c(xC_{kk} + F'') + xC_kC_{kc} = -(e^{-\theta k})(x\theta^2ce^{-\theta k}) + x(-\theta ce^{-\theta k})(-\theta e^{-\theta k}) = -x\theta^2ce^{-2\theta k} + x\theta^2ce^{-2\theta k} = 0$.

³⁹From (28), the effect of a difference in the cost parameter c on the level of investment is given by: $Dk_c = (2p' + xp'')xC_{kc} + C_cC_k$. In general the first term on the right-hand side is ambiguous in sign while the second is negative. In the Spence case, the first term is positive and dominates the second, and the expression as a whole simplifies to: $k_c = \theta c$.

6.4.3 R&D with Threshold Effects

The fact that the specification due to Spence is just on the threshold between super- and submodularity has implausible implications as we have seen. It also suggests that a less convex investment cost function would imply submodularity. Such a specification is found by adapting that of Spence as follows:

$$C(c, k) = c_0 + ce^{-\theta k^a}, a > 0 \quad F(k) = k \quad (32)$$

This is illustrated in Figure 7(b) (drawn for $c_0 = \theta = 1$ and $a = 2$).⁴⁰ For values of a greater than the Spence case of $a = 1$, the cost function is initially concave and then becomes convex.⁴¹ This justifies calling this specification one of threshold effects in R&D: low levels of investment have a relatively small effect on production costs whereas higher levels yield a larger payoff. In the FDI context this implies that firms will select into different modes of market access in exactly the opposite way to Proposition 2. Since profits are submodular in t and c , less efficient firms have a greater incentive to establish a foreign affiliate and carry out their R&D investment locally. By contrast, more efficient firms gain relatively little from further investment in R&D, and find it more profitable to concentrate production in their home plant and serve foreign markets by exporting. Hence the conventional sorting is reversed.

7 Conclusion

This paper has provided a novel approach to the central question of how different firms select into different modes of serving a market. In doing so, we give a unifying perspective on a large and growing literature, identify the critical assumptions which drive many existing results, and develop an approach which can easily be applied to new ones. The key

⁴⁰This is the Gaussian distribution.

⁴¹From (32), $C_{kk} = -\theta a c k^{a-2} e^{-\theta k^a} (a - 1 - \theta a k^a)$. For $0 < a \leq 1$ this is always positive. However, for $a > 1$ it is negative for low k and then becomes positive. The point of inflection occurs where the expression in brackets is zero, which is independent of c and so (for given a and θ) is the same for all firms. In the case illustrated, with $\theta = 1$ and $a = 2$, this occurs at $k = 1/\sqrt{2}$. No firm will produce positive output below the inflection point, since C_{kk} must be positive from the second-order conditions. Note that, while the function is concave at some points and convex at others, it is log-concave everywhere.

criterion we highlight is extremely parsimonious: all that needs to be checked is whether the function giving the maximum profits a firm can earn in a market is supermodular in the firm's own cost parameter, and in a second parameter measuring the marginal cost of accessing the market. We show that this property is necessary and sufficient for the conventional sorting pattern: more efficient firms select into the lower-variable-cost mode of serving the market, whereas less efficient firms select into the higher-variable-cost mode. We first proved this result for a special case where a single monopoly firm chooses between exporting to a foreign market and engaging in foreign direct investment there. We then showed that our approach generalizes to other market structures, both oligopolistic and monopolistically competitive; and to a wide range of other firm choices, including between outsourcing and producing in-house, and between producing with more or less skill- or R&D-intensive techniques.

Our criterion for selection effects is simple both in what it includes and in what it omits. Our approach does not require any special assumptions about the distribution of firm productivities nor about whether countries are symmetric. We are able to dispense with such assumptions because our approach sidesteps the key issue of existence of equilibrium. As Maskin and Roberts (2008) show in a different context, given that an equilibrium exists, its properties can often be established relatively easily.

Since the impact effect of both tariffs and production costs is to lower profits, it is not so surprising that there are many cases where their cross effect is positive, so that supermodularity holds. Nevertheless, the restriction of supermodularity is a non-trivial one, and we have shown that there are many plausible examples where it does not hold. Perhaps most surprisingly, we have shown that, in the canonical case of horizontal FDI where exports incur iceberg transport costs, supermodularity fails to hold at every level of output for most widely-used demand systems, except the CES. Thus if preferences are quadratic or translog, and if selection is observed, it is likely that the most efficient firms will select into exporting rather than FDI. In the case of horizontal FDI, the source of this anomalous result can be traced to the assumption of iceberg transport costs: when higher productive efficiency translates into lower transport costs, the most efficient firms suffer

a lower transport penalty and so will select into exporting rather than FDI. However, in the case of choosing between outsourcing at home versus offshoring to a lower-wage location, our result continues to hold, and implies that for most non-CES preferences more productive firms will select into outsourcing at home, where their greater efficiency offsets the higher wage penalty they incur. We have also identified other plausible cases where supermodularity may fail, such as fixed costs which are higher for more efficient firms, and market-specific investment costs which are subject to threshold effects.

Our results cast the role of fixed costs as determinants of selection effects in a new light. For example, in the choice between FDI and exports, a fixed cost of FDI is essential for a proximity-concentration trade-off to exist. For FDI to be the preferred mode of market access, a firm must be able to afford the additional fixed cost of FDI. However, fixed costs are irrelevant to the determination of which firms are in that category. What matters is the difference-in-differences effect on profits of the marginal costs of production and trade. When supermodularity prevails, a more efficient firm has relatively higher profits in the low-tariff case, but when submodularity holds, the opposite is true. In this paper we first highlighted the implications of this insight for selection into FDI, and then noted that the general point applies to other cases, including selection by more efficient firms into offshoring as in Antràs and Helpman (2004), or into more skill-intensive techniques as in Bustos (2011). There are likely to be still other models which can be illuminated by our approach, and other contexts where the assumption of supermodularity helps to bound comparative statics responses.

8 Appendix

8.1 Proof of Proposition 1

Sufficiency is trivial. To prove necessity, let $t > 0$. If π is not supermodular in t and c , then there exist some c_1 and c_2 such that $c_1 > c_2$ and $\pi(t, c_1) - \pi(t, c_2) < \pi(0, c_1) - \pi(0, c_2)$.

Rearranging terms gives:

$$\pi(0, c_1) - \pi(t, c_1) > \pi(0, c_2) - \pi(t, c_2) \quad (33)$$

Thus $\gamma(t, c_1, f) > \gamma(t, c_0, f)$ for all f ; i.e., the higher-cost firm has a greater incentive to engage in FDI than in exports. To prove that selection will actually occur for some fixed cost, choose f such that $f = \frac{1}{2} [\pi(0, c_1) - \pi(t, c_1) + \pi(0, c_2) - \pi(t, c_2)]$. Then for this f , we have $\gamma(t, c_1, f) = \pi(0, c_1) - \pi(t, c_1) - f \geq 0$ and $\gamma(t, c_2, f) = \pi(0, c_2) - \pi(t, c_2) - f \leq 0$. Hence the higher-cost firm will serve the foreign market via FDI while the lower-cost firm will serve it by exports. Thus, if π is not supermodular in t and c we can always find a fixed cost such that the conventional sorting is reversed. It follows that supermodularity is necessary for the conventional sorting. \square

8.2 Superconvexity

Throughout, we measure the convexity of a function $y = f(x)$ by the unit-free measure $\rho^f \equiv -\frac{xf''}{f'}$. In particular, the convexity of an inverse demand function $p = p(x)$ is $\rho \equiv -\frac{xp''}{p'}$. Our formal definition of superconvexity is as follows:

Definition 2. Define $P \equiv \ln p$ and $X \equiv \ln x$, and write $P = \hat{p}(X)$. Then $p(x)$ is superconvex if and only if $\hat{p}(x)$ is convex; i.e., if and only if $\hat{\rho} \equiv -\frac{X\hat{p}''}{\hat{p}'} \geq 0$.

Some implications of superconvexity are easily established:

Lemma 4. *Superconvexity implies convexity but the converse does not hold.*

Proof. Differentiating $\hat{p}(X)$ yields: $\hat{p}' = \frac{dP}{dX} = \frac{xp'}{p}$; and $\hat{p}'' = \frac{d\hat{p}'}{dX} = x \frac{d\hat{p}'}{dx} = \frac{x}{p^2} [p(xp'' + p') - x(p')^2] = \frac{xp'}{p} \left(\frac{xp''}{p'} + 1 - \frac{xp'}{p} \right)$. Hence $\hat{\rho}$ equals:

$$\hat{\rho} = -\frac{X\hat{p}''}{\hat{p}'} = X \left(\rho - \frac{\varepsilon + 1}{\varepsilon} \right) \quad (34)$$

Clearly, $\hat{\rho} > 0 \Rightarrow \rho > 0$. \square

Lemma 5. *A demand function is superconvex if and only if it is more convex than a CES demand function with the same elasticity.*

Proof. Differentiating the CES inverse demand function $p = \alpha x^{-1/\sigma}$ gives: $p' = -\alpha \frac{1}{\sigma} x^{-(1+\sigma)/\sigma}$; and $p'' = \alpha \frac{\sigma+1}{\sigma^2} x^{-(1+2\sigma)/\sigma}$. Hence we have $\varepsilon^{CES} = \sigma$ and $\rho^{CES} = \frac{\sigma+1}{\sigma}$, from which the result follows by inspection of equation (34). \square

Lemma 6. *A demand function is superconvex if and only if its elasticity is increasing in sales.*

Proof. Differentiating the expression for the elasticity of demand, $\varepsilon(x) = -\frac{p(x)}{xp'(x)}$, yields:

$$\varepsilon_x = -\frac{1}{x} + \frac{p(p' + xp'')}{(xp')^2} = -\frac{1}{x} (1 + \varepsilon - \varepsilon\rho) \quad (35)$$

Comparison with (34) shows that ε_x and $\hat{\rho}$ always have the same sign: $\varepsilon_x = \frac{\varepsilon}{xX} \hat{\rho}$. \square

Finally, super-convexity can also be expressed in terms of the direct demand function $x = x(p)$, with elasticity $e(p) \equiv -\frac{px'}{x} = \varepsilon[x(p)]$:

Lemma 7. *A demand function is superconvex if and only if its elasticity is decreasing in price.*

Proof. Differentiating the identity equating the two expressions for the elasticity of demand, $e(p) = \varepsilon[x(p)]$, yields $e_p = \varepsilon_x p'$, which from Lemma 6 implies that e_p and $\hat{\rho}$ always have opposite signs: $e_p = p' \frac{\varepsilon}{xX} \hat{\rho}$. \square

8.3 Proof of Proposition 2

Differentiating the profit function $\pi(t, c) = \max_x [p(x) - \tau c] x$ gives: $\pi_t = -cx$; and

$$\pi_{tc} = -x - c \frac{dx}{dc} = -x - \frac{\tau c}{2p' + xp''} \quad (36)$$

We want to express the right-hand side in terms of ε and ε_x . First, solve (35) for $p' + xp''$ in terms of ε_x , and add p' to it. Next, use the definition of ε to eliminate p' , $p' = -\frac{p}{x\varepsilon}$, which gives the second-order condition in terms of ε and ε_x :

$$2p' + xp'' = -\frac{p}{x\varepsilon^2} (\varepsilon - 1 - x\varepsilon_x) \quad (37)$$

This confirms that the second-order condition $2p' + xp'' < 0$ is equivalent to $\varepsilon - 1 - x\varepsilon_x > 0$. The last preliminary step is to use the first-order condition $p - \tau c + xp' = 0$ to express τc in terms of p and ε : $\tau c = p + xp' = p - \frac{p}{\varepsilon} = \frac{\varepsilon - 1}{\varepsilon}p$. (This is very familiar in the CES case.) Finally, substitute these results into (36):

$$\pi_{tc} = -x + \frac{\varepsilon - 1}{\varepsilon - 1 - x\varepsilon_x} \varepsilon x \quad (38)$$

Collecting terms gives the desired expression in (15). \square

8.4 Proof of Proposition 3

The proof follows immediately by substituting for ε_x from (35) into the expression for π_{tc} in (15), and noting that, as in (15), the denominator must be positive from the second-order condition.

8.5 Proof of Lemma 2

For the demand function $p = (x - \beta)^{-1/\sigma}$, we have $p' = -\frac{1}{\sigma}(x - \beta)^{-\frac{\sigma+1}{\sigma}}$ and $p'' = \frac{\sigma+1}{\sigma^2}(x - \beta)^{-\frac{2\sigma+1}{\sigma}}$. Hence $\varepsilon = \frac{x-\beta}{x}\sigma$ and $\rho = \frac{x}{x-\beta}\frac{\sigma+1}{\sigma}$. It follows immediately that $\varepsilon_x = \frac{\beta}{x^2}\sigma$, and so the demand function is strictly subconvex ($\varepsilon_x < 0$) if and only if β is negative. To establish for which values of σ the profit function is supermodular, rewrite the elasticity as $\varepsilon = \frac{\sigma+1}{\rho}$. Hence we seek to minimize $\sigma = \varepsilon\rho - 1$ by choice of ε and ρ , subject to the supermodularity constraint from Proposition 3, $\varepsilon + \rho \geq 3$. Solving gives the boundary values $\varepsilon^* = \rho^* = 1.5$, which imply that the threshold value of σ is $\sigma^* = \left(\frac{3}{2}\right)^2 - 1 = 1.25$. \square

8.6 Proof of Lemma 3

Maximizing operating profits, $\pi = (p - \tau c)x$, yields the first-order condition, $A - 2bx = \tau c$, which can be solved for optimal output: $x = \frac{1}{2b}(A - \tau c)$. Substituting back into the expression for profits gives the maximized operating profit function:

$$\pi(t, c) = bx^2 = \frac{1}{4b}(A - \tau c)^2 \quad (39)$$

Hence the second cross-derivative is:

$$\pi_{tc} = -x + \frac{\tau c}{2b} = -\frac{1}{2b} (A - 2\tau c) \quad (40)$$

This is clearly positive for $c \geq \frac{A}{2\tau}$ and negative for $c \leq \frac{A}{2\tau}$, which proves the Lemma. \square

References

- ANTRÀS, P., AND E. HELPMAN (2004): “Global Sourcing,” *Journal of Political Economy*, 112(3), 552–580.
- ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2009): “New trade models, same old gains?,” NBER Working Paper No. 15628, December; forthcoming in *American Economic Review*.
- ATHEY, S. (2002): “Monotone Comparative Statics Under Uncertainty,” *Quarterly Journal of Economics*, 117(1), 187–223.
- BEHRENS, K., G. MION, AND G. I. OTTAVIANO (2010): “Economic Integration and Industry Reallocations: Some Theory with Numbers,” forthcoming in M.N. Jovanovic (ed.): *International Handbook of Economic Integration*, Cheltenham, UK: Edward Elgar.
- 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.
- COSTINOT, A. (2007): “Heterogeneity and Trade,” University of California at San Diego, Economics Working Paper Series.
- (2009): “An Elementary Theory of Comparative Advantage,” *Econometrica*, 77(4), 1165–1192.
- COSTINOT, A., AND J. VOGEL (2010): “Matching and Inequality in the World Economy,” *Journal of Political Economy*, 118(4), 747–786.
- D’ASPREMONT, C., AND A. JACQUEMIN (1988): “Cooperative and Noncooperative R&D in Duopoly with Spillovers,” *American Economic Review*, 78(5), 1133–1137.
- DAVIES, R. B., AND C. ECKEL (2010): “Tax Competition for Heterogeneous Firms with Endogenous Entry,” *American Economic Journal: Economic Policy*, 2(1), 77–102.

- DHINGRA, S., AND J. MORROW (2011): “The Impact of Integration on Productivity and Welfare Distortions Under Monopolistic Competition,” draft, Princeton University.
- EGGER, H., AND U. KREICKEMEIER (2009): “Firm Heterogeneity and the Labor Market Effects of Trade Liberalization,” *International Economic Review*, 50(1), 187–216.
- FEENSTRA, R. C. (2003): “A Homothetic Utility Function for Monopolistic Competition Models, without Constant Price Elasticity,” *Economics Letters*, 78(1), 79–86.
- FORSLID, R., T. OKUBO, AND K. H. ULLTVEIT-MOE (2011): “International Trade, CO2 Emissions and Heterogeneous Firms,” CEPR Discussion Paper DP8583, Centre for Economic Policy Research.
- GHIRONI, F., AND M. J. MELITZ (2005): “International Trade and Macroeconomic Dynamics with Heterogeneous Firms,” *Quarterly Journal of Economics*, 120(3), 865–915.
- GROSSMAN, G. M., AND E. HELPMAN (2002): “Integration Versus Outsourcing in Industry Equilibrium,” *Quarterly Journal of Economics*, 117(1), 85–120.
- GROSSMAN, G. M., AND G. MAGGI (2000): “Diversity and Trade,” *American Economic Review*, 90(5), 1255–1275.
- GROSSMAN, S., AND O. HART (1986): “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *Journal of Political Economy*, pp. 691–719.
- HART, O., AND J. MOORE (1990): “Property Rights and the Nature of the Firm,” *Journal of Political Economy*, pp. 1119–1158.
- HELPMAN, E., O. ITSKHOKI, AND S. REDDING (2010): “Inequality and Unemployment in a Global Economy,” *Econometrica*, 78(4), 1239–1283.
- HELPMAN, E., M. J. MELITZ, AND S. R. YEAPLE (2004): “Export Versus FDI with Heterogeneous Firms,” *American Economic Review*, 94(1), 300–316.

- KINGMAN, J. F. C. (1961): “A Convexity Property of Positive Matrices,” *Quarterly Journal of Mathematics*, 12(1), 283–284.
- KLEINERT, J., AND F. TOUBAL (2006): “Distance Costs and Multinationals’ Foreign Activities,” Centre for Economic Institutions Working Paper No. 2006-6, Hitotsubashi University.
- (2010): “Gravity for FDI,” *Review of International Economics*, 18(1), 1–13.
- KRUGMAN, P. R. (1979): “Increasing Returns, Monopolistic Competition, and International Trade,” *Journal of International Economics*, 9(4), 469–479.
- LEAHY, D., AND C. MONTAGNA (2009): “Outsourcing vs FDI in Oligopoly Equilibrium,” *Spatial Economic Analysis*, 4(2), 149–166.
- LIMÃO, N. (2005): “Trade Policy, Cross-Border Externalities and Lobbies: Do Linked Agreements Enforce More Cooperative Outcomes?,” *Journal of International Economics*, 67(1), 175–199.
- MASKIN, E. S., AND K. W. ROBERTS (2008): “On the Fundamental Theorems of General Equilibrium,” *Economic Theory*, 35(2), 233–240.
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71(6), 1695–1725.
- MELITZ, M. J., AND G. I. OTTAVIANO (2008): “Market Size, Trade, and Productivity,” *Review of Economic Studies*, 75(1), 295–316.
- MILGROM, P., AND J. ROBERTS (1990): “The Economics of Modern Manufacturing: Technology, Strategy, and Organization,” *American Economic Review*, 80(3), 511–528.
- MILGROM, P., AND C. SHANNON (1994): “Monotone Comparative Statics,” *Econometrica*, pp. 157–180.
- MRÁZOVÁ, M., AND J. P. NEARY (2010): “Firm Selection into Export-Platform Foreign Direct Investment,” presented to the ETSG Annual Conference, Lausanne.

- NEARY, J. P. (2002): “Foreign Direct Investment and the Single Market,” *Manchester School*, 70(3), 291–314.
- (2009): “Putting the ‘New’ into New Trade Theory: Paul Krugman’s Nobel Memorial Prize in Economics,” *Scandinavian Journal of Economics*, 111(2), 217–250.
- (2010): “Two and a Half Theories of Trade,” *The World Economy*, 33(1), 1–19.
- NEFUSSI, B. (2006): “Exports versus FDI: Evidence from Two French Industries,” Working paper, Crest-INSEE.
- OLDENSKI, L. (2010): “Export Versus FDI: A Task-Based Approach,” mimeo., Georgetown University.
- PETIT, M.-L. P., AND F. SANNA-RANDACCIO (2000): “Endogenous R&D and Foreign Direct Investment in International Oligopolies,” *International Journal of Industrial Organization*, 18(2), 339–367.
- PORTER, L. A. (2011): “Asymmetric Oligopoly and Foreign Direct Investment: Implications for Host-Country Tax-Setting,” forthcoming in *International Economic Journal*.
- SPENCE, M. (1984): “Cost Reduction, Competition, and Industry Performance,” *Econometrica*, 52(1), 101–122.
- WILLIAMSON, O. (1975): *Markets and Hierarchies: Analysis and Antitrust Implications: A Study in the Economics of Internal Organization*. New York: Free Press.
- ZHELOBODKO, E., S. KOKOVIN, M. PARENTI, AND J.-F. THISSE (2011): “Monopolistic Competition: Beyond the CES,” Working Paper 2011 - 08, Paris School of Economics.