Heterogeneous firm Models

[*My notes are in beta. If you see something that doesn't look right, I would greatly appreciate a heads-up.*]

Two facts about exporters:

1. Not many firms export

2. Firms that do export are larger and more productive

As firm-level data became available, these two facts suggested that heterogeneity could be important for understanding international economics. A potential explanation for these facts is that firms self-select into exporting. This mechanism, in various guises, shows up throughout international trade and macro models.

The big idea: Firms pay a fixed cost to sell into a market. Firms are heterogeneous in the profit they earn from selling into that market. Thus, more profitable firms select into serving the market.

1   Homogeneous firm trade models

Before turning to heterogenous firm models, it is useful to see how the models work with symmetric firms. The classic reference here is Krugman (1980). [Todo: Write this up]

2   Heterogeneous firm trade models

The analysis here mostly follows Chaney (2008), which makes one significant departure from the Melitz (2003) framework: In the Melitz framework the mass of entrants (M in Melitz's notation) is endogenous and potential firms pay a cost f_e to draw from the productivity distribution. This is essentially a static version of Hopenhayn (1992). In Chaney's framework, the mass of potential entrants is fixed, so there are aggregate profits in equilibrium.

Endowments and preferences
There are N countries, each endowed with L_i units of labor and μ_i potential firms. By potential firms, we mean that country i gets μ_i draws from the productivity distribution.

In each country, two types of goods are produced: a homogeneous good whose consumption in i is denoted c_{0i} and a continuum of differentiated goods, indexed by ω. The consumer inelastically supplies labor to the firms, and, as the owner of the firms, receives the firms' profits. The consumer in country i chooses consumption of the homogenous and differentiated goods to maximize

$$(1-\alpha)\log c_{0i} + \frac{\alpha\sigma}{\sigma-1}\log\int_0^m c_i(\omega)^{\frac{\sigma-1}{\sigma}}\,d\omega \tag{1}$$

subject to the consumer's budget constraint

$$p_{0i}c_{0i} + \int_0^m p_i(\omega)c_i(\omega)\,d\omega = w_iL_i + \pi_i, \tag{2}$$

and the usual non-negativity constraints. The wage is denoted w_i and the profit from all firms in operation in country i is π_i. (Chaney makes a different assumption about firm ownership. We discuss this later.) The total mass of differentiated goods consumed is m. We will later make

clear the relationship between the mass of firms in operation, the mass of potential firms, and the probability distribution over potential firms. The solution to the consumer's problem leads to the well known CES demand functions for differentiated goods

$$c_i(\omega) = \alpha \left(w_i L_i + \pi_i \right) p_i(\omega)^{-\sigma} P_i^{\sigma-1}, \tag{3}$$

where the aggregate price index for differentiated goods is

$$P_i = \left(\int_0^m p_i(\omega)^{-(\sigma-1)} \, d\omega \right)^{\frac{-1}{\sigma-1}}. \tag{4}$$

Firms

The homogenous consumption good is produced using a constant returns to scale production function that is identical across countries, $y_0 = \ell_0$, and sold in competitive markets. The homogenous good is freely traded, so we take it to be the numeraire good and set $p_{0i} = w_i = 1$ for all i. This is a common trick: The homogeneous good pins down the wage, so it effectively makes the model a partial equilibrium model. This makes solving the model analytically much simpler, but this is not necessary (and probably not desirable) if you are willing to turn on a computer to solve the model.

Differentiated-good firms are monopolistic competitors. The differentiated-good firm in market i producing good ω for sale in market j has the production function $y_{ij}(\omega) = \varphi(\omega)\ell_{ij}(\omega)$[1]. The firm producing ω has marginal productivity $\varphi(\omega) \in [\underline{\varphi}, \infty)$. The fixed operating cost of selling to market j is κ_{ij}, which is paid in units of labor. This is a static model, so there is no sense in which this is a "sunk cost" as in Roberts and Tybout (1997). When we move to dynamic models, the concept of a sunk cost will become meaningful.

Firms produce with a constant marginal cost, (gross of fixed costs) so the maximization problem of a firm can be separated across destination markets. The firm takes as given the residual demand for its good and chooses its price to maximize profits

$$\max_{p_{ij}, \ell_{ij}} \pi_{ij}(\omega) = p_{ij}(\omega)c_{ij}(\omega) - w_i \ell_{ij}(\omega) - w_i \kappa_{ij} \tag{5}$$

$$\text{s.t. } \tau_{ij} c_{ij}(\omega) = \varphi(\omega)\ell_{ij}$$

which yields the usual constant markup over marginal pricing condition

$$p_{ij}(\omega) = \frac{\sigma}{\sigma - 1} \frac{w_i}{\varphi(\omega)} \tau_{ij}. \tag{6}$$

The first term is the markup, which is a function of the elasticity of substitution. When goods are closer substitutes, consumers are more sensitive to prices, so markups are lower. The second term is the firm's marginal cost and the last term is the adjustment for variable trade costs. Notice that prices do not differ across markets except to compensate for variable trade costs. This is not a model in which firms price discriminate across markets. (Although there is evidence in the data

---

[1]If we write the production function as $y_{ij}(\omega) = \varphi(\omega)^{1/(\sigma-1)}\ell_{ij}(\omega)$, the variance of the firm-size distribution will not be dependent on σ. To see this, solve for labor demand and notice that it is proportional to φ in this case.

that they do.)

Notice that the price set by a firm (and in fact, any other decision made by the firm) depends not on the good's name, ω, but on its productivity φ. Since all firms with productivity φ will behave identically, and varieties enter utility symmetrically, we no longer characterize a good by its label, but by its productivity. (We did this in the Ricardian models, too.)

A firm's sales (which are exports when $i \neq j$) are

$$p_{ij}(\varphi)c_{ij}(\varphi) = \alpha \left(w_j L_j + \pi_j\right) P_j^{\sigma-1} \left(\frac{\sigma}{\sigma-1} \frac{w_i \tau_{ij}}{\varphi}\right)^{-(\sigma-1)} \tag{7}$$

and profits are a constant fraction of firm revenues (another special property of CES demand + monopolistic competition)

$$\pi_{ij}(\varphi) = \sigma^{-1}\alpha \left(w_j L_j + \Pi_j\right) P_j^{\sigma-1} \left(\frac{\sigma}{\sigma-1} \frac{w_i \tau_{ij}}{\varphi}\right)^{-(\sigma-1)} - w_i \kappa_{ij}. \tag{8}$$

We will characterize an equilibrium in which there exists, for each country pair, a firm with productivity $\hat{\varphi}_{ij}$ that is indifferent to selling in market j. We will assume that the fixed costs are large enough that this cutoff productivity level is always greater than the minimum productivity level $\hat{\varphi}_{ij} > \underline{\varphi}$. This cutoff level of productivity is defined as the level of productivity that yields a firm with productivity $\hat{\varphi}_{ij}$ zero profits from selling in j, $\pi_{ij}(\hat{\varphi}_{ij}) = 0$,

$$\hat{\varphi}_{ij} = \frac{\sigma}{\sigma-1} \frac{w_i \tau_{ij}}{P_j} [\sigma^{-1}\alpha Y_j]^{-\frac{1}{\sigma-1}} (w_i \kappa_{ij})^{\frac{1}{\sigma-1}}, \tag{9}$$

where $Y_j = w_j L_j + \pi_j$ is total income in country j. Since profits are increasing in productivity, all firms with $\varphi \geq \hat{\varphi}_{ij}$ will sell into market j. This is the key mechanism in the model: The fixed costs induce selection over firm productivity. The more productive firms can earn large enough profits to cover fixed costs, so they will export to market j. This is one of the stylized data facts: Not all firms export.

More productive firms sell more than less productive firms (see equation 7) because a firm prices on the elastic part of its demand curve. Exporters will be larger than nonexporters both in terms of total sales, domestic sales, and employment. This is another stylized data fact: Exporters are large.

If you can grok this figure, then you understand 95 percent of how this model works. The other 5 percent is general equilibrium stuff.

Firm heterogeneity
Let the distribution over firm types be Pareto

$$G(x) = 1 - \underline{x}^{\gamma} x^{-\gamma}. \tag{10}$$

The parameter γ governs the amount of variance in the distribution. We need $\gamma > 2$ for the variance to be finite. We will also require that $\gamma > \sigma - 1$. Technically, we need this to get some integrals to converge, but, intuitively, this restriction says that, relative to the variance in firm productivity, we can't have too much substitution. If goods are too substitutable, then consumers

Figure 1: Cutoff productivity.
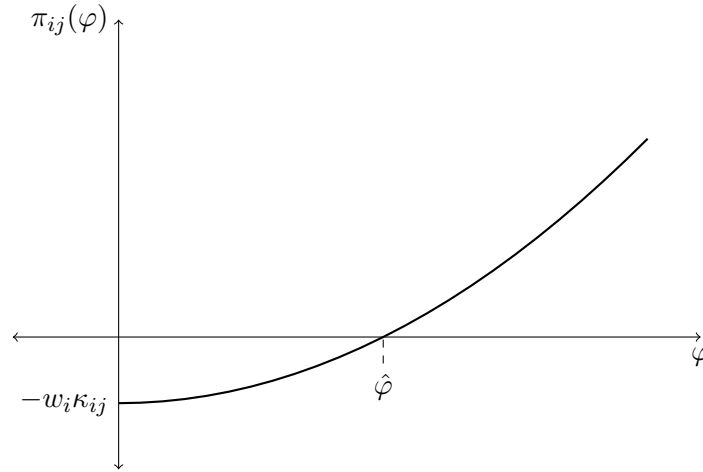
will want to push all of their spending into the low-priced (high productivity) firms, and the model explodes. This restriction keeps consumers demanding "enough" of the higher priced goods.

The Pareto distribution has the special property that truncating it from below creates a distribution that is still Pareto with shape parameter γ, but with a different lower bound, $\underline{x}'$. This property will be useful, given that these kinds of models are characterized by truncations of the productivity distribution. (This is analogous to the Eaton and Kortum (2002) + Frechet assumption.)

In equilibrium, the firm size distribution in this model will be something close to the underlying distribution we assume over productivity. (This is not a very deep model of the firm size distribution.) The U.S. firm size distribution is roughly Pareto, at least in the right tail (see Axtell, 2001), although there is some disagreement about this.

Aggregates
With the model characterized by the cutoff values, we can now work out some aggregates. The aggregate price index is

$$P_j^{-(\sigma-1)} = \sum_{k=1}^{N} \int_{\hat{\varphi}_{kj}} (w_k \tau_{kj})^{-(\sigma-1)} ((\sigma-1)/\sigma)^{\sigma-1} \varphi^{\sigma-1} \mu_k \gamma \underline{\varphi}^{\gamma} \varphi^{-\gamma-1} d\varphi \tag{11}$$

$$P_j^{-(\sigma-1)} = \sum_{k=1}^{N} \frac{(w_k \tau_{kj})^{-(\sigma-1)} ((\sigma-1)/\sigma)^{\sigma-1} \mu_k \gamma \underline{\varphi}^{\gamma}}{\gamma - (\sigma-1)} \hat{\varphi}_{kj}^{\sigma-1-\gamma} \tag{12}$$

$$P_j^{-(\sigma-1)} = \sum_{k=1}^{N} \frac{(w_k \tau_{kj})^{-(\sigma-1)} ((\sigma-1)/\sigma)^{\sigma-1} \mu_k \gamma \underline{\varphi}^{\gamma}}{\gamma - (\sigma-1)} \left(\frac{w_k \tau_{kj} \sigma}{(\sigma-1)P_j} \right)^{(\sigma-1)-\gamma} [\sigma^{-1} \alpha Y_j]^{\gamma \frac{1-\rho}{\rho}-1} \kappa_{kj}^{1-\gamma \frac{1-\rho}{\rho}} \tag{13}$$

$$P_j = \lambda_2 Y_j^{\frac{1}{\gamma}-\frac{1}{\sigma-1}} \Theta_j \tag{14}$$

where λ_2 is a constant, $\lambda_2 = [\sigma^{-1}\alpha]^{\frac{1}{\gamma}-\frac{1}{\sigma-1}}\frac{\sigma}{\sigma-1}\left(\frac{\gamma\varphi^\gamma}{\gamma-(\sigma-1)}\right)^{-1/\gamma}$, and

$$\Theta_j^{-\gamma} = \sum_{k=1}^{N} \mu_k(w_k\tau_{kj})^{-\gamma}\kappa_{kj}^{1-\gamma\frac{1}{\sigma-1}}. \tag{15}$$

Notice that Θ has the the potential number of firms (μ), the cost of inputs (w), and the costs of trading (τ, κ) for all the other countries selling into j. This a multilateral resistance kind of term.

Inspecting (13), we see that bigger markets ($1/\gamma - 1/(\sigma-1) < 0$) have lower prices. Bigger markets allow more firms to cover the fixed cost of serving j, so there are more varieties.

Using (14), we can rewrite the cutoff productivity as

$$\hat{\varphi}_{ij} = \lambda_4 Y_j^{-1/\gamma}\frac{w_i\tau_{ij}}{\Theta_j}\kappa_{ij}^{\frac{1}{\sigma-1}} \tag{16}$$

$$\hat{\varphi}_{ij}^{-\gamma} = \lambda_4^{-\gamma}Y_k\left(\frac{w_i\tau_{ij}}{\Theta_j}\right)^{-\gamma}\kappa_{ij}^{-\gamma\frac{1}{\sigma-1}} \tag{17}$$

$$\hat{\varphi}_{ij}^{\sigma-1-\gamma} = \lambda_4^{\sigma-1-\gamma}Y_j^{1-\frac{\sigma-1}{\gamma}}\left(\frac{w_i\tau_{ij}}{\Theta_j}\right)^{\sigma-1-\gamma}\kappa_{ij}^{1-\frac{\gamma}{\sigma-1}} \tag{18}$$

where $\lambda_4 = [\sigma^{-1}\alpha]^{\frac{-1}{\gamma}}\left(\frac{\gamma\varphi^{-\gamma}}{\gamma-(\sigma-1)}\right)^{\frac{1}{\gamma}}$. We will use $\hat{\varphi}_{ij}^{-\gamma}$ and $\hat{\varphi}_{ij}^{\sigma-1-\gamma}$ in later derivations, so it's useful to write them out here.

Aggregate profits in country i are the sum of the profits of all firms owned by country i that are active across each market, k. They are made up of two parts: the variable profits earned by firms, minus the fixed costs paid by firms.

$$\pi_i = \sum_{k=1}^{N}\int_{\hat{\varphi}_{ik}}\sigma^{-1}\alpha Y_k\left(\frac{\sigma}{\sigma-1}\frac{w_i\tau_{ik}}{P_k}\right)^{-(\sigma-1)}\varphi^{\sigma-1}\mu_i\gamma\underline{\varphi}^\gamma\varphi^{-\gamma-1}d\varphi - \sum_{k=1}^{N}\kappa_{ik}\mu_i\underline{\varphi}^\gamma\hat{\varphi}_{ik}^{-\gamma} \tag{19}$$

The first term is

$$\pi_i^{(1)} = \alpha\sigma^{-1}\mu_i\sum_{k=1}^{N}Y_k\left(\frac{w_i\tau_{ik}}{\Theta_k}\right)^{-\gamma}\kappa_{ik}^{1-\frac{\gamma}{\sigma-1}}, \tag{20}$$

the second term is

$$\pi_i^{(2)} = \sum_{k=1}^{N}\kappa_{ik}\mu_i\underline{\varphi}^\gamma\hat{\varphi}_{ik}^{-\gamma} = \mu_i\underline{\varphi}^\gamma\lambda_4^{-\gamma}\sum_{k=1}^{N}Y_k\left(\frac{w_i\tau_{ik}}{\Theta_k}\right)^{-\gamma}\kappa_{ik}^{1-\gamma\frac{1}{\sigma-1}}, \tag{21}$$

and aggregate profits are

$$\pi_i = \frac{\alpha(\sigma-1)}{\gamma\sigma}\sum_{k=1}^{N}Y_k\frac{\mu_i(w_i\tau_{ik})^{-\gamma}\kappa_{ik}^{1-\gamma\frac{1-\rho}{\rho}}}{\Theta_k^{-\gamma}}. \tag{22}$$

Let $\theta_{ik} = \mu_i(w_i\tau_{ik})^{-\gamma}\kappa_{ik}^{1-\gamma\frac{1}{\sigma-1}}$, so that $\sum_{i=1}^{N}\theta_{ik} = \Theta_k^{-\gamma}$. [To do: Clean up the definition of Θ to get

rid of that power.] Aggregate profits form a system of equations,

$$\pi_i = \frac{\alpha(\sigma-1)}{\sigma\gamma} \sum_{k=1}^{N} (w_k L_k + \pi_k) \frac{\theta_{ik}}{\Theta_k^{-\gamma}} \quad i = 1, \ldots, N. \tag{23}$$

Note that this is a system of N linear equations in N unknowns — the π_i. In a perfectly symmetric world, we would have $\pi_i = \pi = \alpha(\sigma-1)/(\sigma\gamma - \alpha(\sigma-1))$. Solve this to find the equilibrium of the model.

Existence of equilibrium

Let's set up the system defined in (23) as a matrix equation. Rearrange terms to yield

$$\pi_i - \frac{\alpha(\sigma-1)}{\gamma\sigma} \sum_{k=1}^{N} \pi_k \frac{\theta_{ik}}{\Theta_k^{-\gamma}} = \frac{\alpha(\sigma-1)}{\gamma\sigma} \sum_{k=1}^{N} \frac{\theta_{ik}}{\Theta_k^{-\gamma}} w_k L_k \quad i = 1, \ldots, N. \tag{24}$$

If we define

$$b_i = \frac{\alpha(\sigma-1)}{\gamma\sigma} \sum_{k=1}^{N} \frac{\theta_{ik}}{\Theta_k^{-\gamma}} w_k L_k \quad \text{and} \quad d_{ik} = \frac{\alpha(\sigma-1)}{\gamma\sigma} \sum_{k=1}^{N} \frac{\theta_{ik}}{\Theta_k^{-\gamma}}, \tag{25}$$

we can write the system in (23) as

$$[I - D]\pi = b \tag{26}$$

which facilitates a simple proof of existence: The column sums of D are $\alpha(\sigma-1)/(\gamma\sigma) < 1$ and the vector b is strictly positive, so the Perron-Frobenius theorem tells us that the vector π which solves (26) is strictly positive. In fact, $[I-D]^{-1}$ is strictly positive. [To do: There is more that we can show.]

Chaney's financial structure
In our derivation, we have assumed that each country's agents own the firms that are created in the country. Chaney (2008) assumes that $Y_i = w_i L_i + w_i L_i \pi$, which he interprets as each country owning shares in a global mutual fund which has a dividend of $\pi = \sum_j \pi_j / \sum_j w_j L_j$, and each country owns shares in proportion to its labor income.

$$\pi_i = \frac{\alpha(\sigma-1)}{\gamma\sigma} \sum_{k=1}^{N} (w_k L_k + w_k L_k \pi) \frac{\theta_{ik}}{\Theta_k^{-\gamma}} \tag{27}$$

$$\pi_i = \frac{\alpha(\sigma-1)}{\gamma\sigma} (1+\pi) \sum_{k=1}^{N} w_k L_k \frac{\theta_{ik}}{\Theta_k^{-\gamma}} \tag{28}$$

Take the summation over i,

$$\sum_{i=1}^{N} \pi_i = \frac{\alpha(\sigma-1)}{\gamma\sigma} (1+\pi) \sum_{i=1}^{N} \sum_{k=1}^{N} w_k L_k \frac{\theta_{ik}}{\Theta_k^{-\gamma}}, \tag{29}$$

and note that both the double summation on the right-hand side is equal to world labor income,

$$\sum_{i=1}^{N} \pi_i = \frac{\alpha(\sigma-1)}{\gamma\sigma}(1+\pi)\sum_{i=1}^{N} w_k L_k \tag{30}$$

$$\pi = \frac{\alpha(\sigma-1)}{\gamma\sigma}(1+\pi). \tag{31}$$

In Chaney's model, $\pi = \alpha(\sigma-1)/(\sigma\gamma - \alpha(\sigma-1))$, and $Y_i = w_i L_i \gamma\sigma/(\sigma\gamma - \alpha(\sigma-1))$. Notice that Chaney's assumption yields — even in a model with asymmetric trade costs — the same country-level profits as in the general model with symmetric countries.

Trade flows

Aggregate trade flows from i to j

$$X_{ij} = \int_{\hat{\varphi}_{ij}} \alpha \left(w_j L_j + \pi_j\right) P_j^{\sigma-1} \left(\frac{\sigma}{\sigma-1}\frac{w_i \tau_{ij}}{\varphi}\right)^{-(\sigma-1)} \mu_i \gamma \underline{\varphi}^{\gamma}\varphi^{-\gamma-1}d\varphi \tag{32}$$

$$X_{ij} = \alpha Y_j \frac{\mu_i \left(w_i \tau_{ij}\right)^{-\gamma} \kappa_{ij}^{-\left(\frac{\gamma}{\sigma-1}-1\right)}}{\Theta_j^{-\gamma}} = \alpha Y_j \frac{\theta_{ij}}{\Theta_j^{-\gamma}}. \tag{33}$$

Aggregate trade is a share of the total expenditure on heterogenous goods in country j, where that share is determined by the mass of potential firms, input costs, variable costs, and fixed costs in all of the countries, relative to those in country i.

Gravity and the intensive and extensive margins

This style of model has both an intensive and extensive margin of trade. The extensive margin refers to a firm exporting anything — it is a binary concept. The intensive margin refers to how much a firm exports conditional on exporting. Understanding how these two margins are affected by trade barriers is key to understanding how total trade behaves.

In Chaney's model, μ_i is proportional to L_i, so with a little manipulation, gravity is [To do: derive this from above.]

$$X_{ij} = \alpha \times \frac{Y_i \times Y_j}{Y} \times \left(\frac{w_i \tau_{ij}}{\Theta_j}\right)^{-\gamma} \times (\kappa_{ij})^{-(\gamma\frac{1}{\sigma-1}-1)} \tag{34}$$

1. Again, unit income elasticities.

2. The Θ_j terms play the role of a multilateral resistance. Not only do the costs of production in the exporting country, bilateral trade costs, and the bilateral fixed cost matter, but the costs of production, bilateral trade costs, and fixed costs of all other countries matter, too.

3. The elasticity with respect to variable trade costs is γ and not $\sigma - 1$. This is similar to the Eaton and Kortum (2002) result that a parameter that controls the amount of heterogeneity in producers is key, not the elasticity of substitution. As in Eaton and Kortum, this is the result of the a special choice in distribution. In this case, it is the memoryless property of the Pareto distribution.

   The importance of γ has a nice interpretation. Higher γ implies more variance in the distribution, so there is more potential for competition. This result goes away if we move from

the Pareto distribution to, say, the log normal distribution.

4. Since $\gamma > \sigma - 1$, the elasticity of trade is greater than it is in a model without heterogeneity.

5. The elasticity of trade with respect to fixed costs is $\gamma/(\sigma - 1) - 1$. Lower elasticities in consumption increase the response to a change in entry barriers. When the elasticity of substitution is low, the impact of trade barriers is greater. This is due to the extensive margin. When goods have low substitutability, then the profit a monopolistic competitor can make is higher. Thus, a change in trade barriers has a larger increase in export profit, so the extensive margin response is larger.

Figure 2: Decrease in variable (left panel) and fixed (right panel) trade costs.
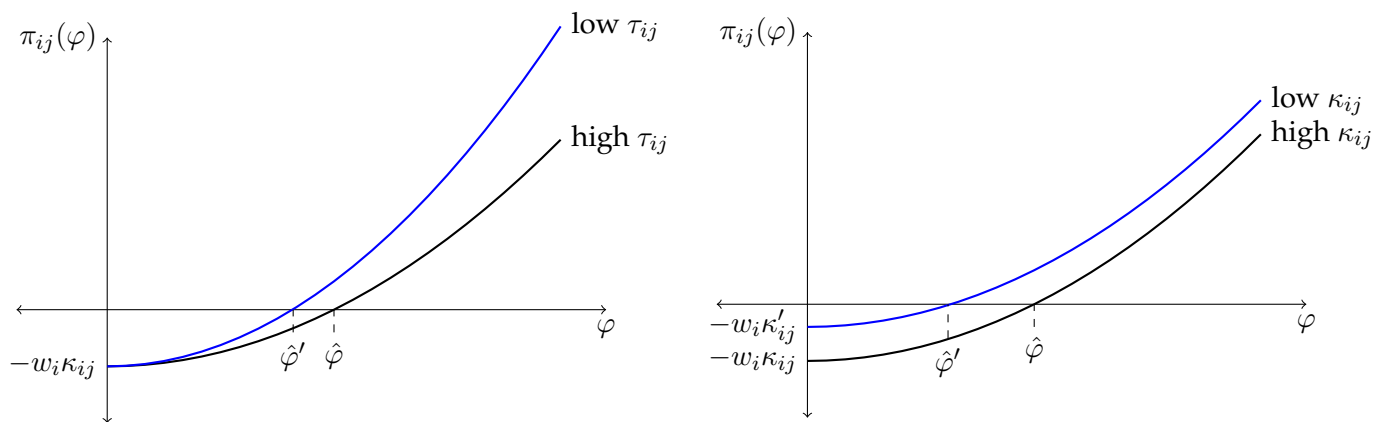
What is lurking below results #3 and #5? When we lower the variable trade costs, two things happen.

1. There is an intensive margin response. Firms that were already exporting now export more. This follows directly from (7), and notice that the elasticity in the intensive margin is $\sigma - 1$.

2. There is an extensive margin response. In the left panel of figure 2, we can see that decreasing the variable trade costs increases the profits a firm can earn which shifts the break even productivity to the left. More firms are now exporting.

Completely differentiate (34) with respect to variable costs and find that

$$-\frac{d \log X_{ij}}{d \log \tau_{ij}} = (\sigma - 1) + (\gamma - (\sigma - 1)) = \gamma \tag{35}$$

The first term on the right hand side is the intensive margin contribution to the total elasticity, and it is related to substitutability in preferences — just like we would expect. The second term is the elasticity from the extensive margin. It depends on both the variance in the productivity distribution and the elasticity of substitution, and it exactly cancels out the intensive margin effect. (Again, this is a Pareto distribution result.)

When we lower the fixed cost of exporting the intensive margin effect is generally small. Once a firm has decided to enter a market and pay the fixed cost, the fixed cost no longer matters for the firm's decisions. The firm's price and sales volume do not depend on the fixed cost. There may

be some impact on the intensive margin through the general equilibrium effects on w_i and π_i, but the Chaney (2008) model kills off these effects. (The effects would likely be small, anyway.) If we Completely differentiate (34) with respect to fixed costs, we have

$$-\frac{d \log X_{ij}}{d \log \kappa_{ij}} = 0 + \frac{\gamma}{\sigma - 1} - 1. \tag{36}$$

Here, there is no intensive margin response. This is because we have effectively killed of the general equilibrium response of wages and incomes to a change in trade barriers. This is all extensive margin, and the intuition follows from #5, above.

Computation

The Pareto versions of the model can be solved directly by solving the system of linear equations in (23), but it is not very hard to solve the more general model.

In a model with the outside good, the wages are pinned down, and the algorithm below is enough to compute an equilibrium. In a model without an outside good, wages are endogenous. The algorithm below is then an "inner loop" and an outer loop can be added to also solve for equilibrium wages.

1. Guess π_i and P_i for all i (this is $2 \times N$ unknowns)

2. Use these guesses to compute cutoff productivities $\hat{\varphi}_{ij}$

3. Using the $\hat{\varphi}_{ij}$, compute the implied aggregate price indices using (11). Let's call them P_i'. This requires computing an integral. If you have not done this before, Google "quadrature."

4. Using the $\hat{\varphi}_{ij}$ and the guesses from step 1, compute the implied aggregate profits using (19). Let's call them π_i'. This also requires computing an integral.

5. Check: Does $\pi_i' = \pi_i$ and $P_i' = P_i$ for all i? If not, go back to step 1 and update the guesses.

This algorithm defines a system of $2N$ nonlinear equations in $2N$ unknowns. Write a function that takes the P_i and π_i as arguments and returns the differences defined in step 5. Pass this function to a nonlinear equation solver, like fsolve in MATLAB — done!

To solve a model with a free entry condition and an endogenous mass of entrants, set $\pi_i = 0$, and choose μ_i (the mass of potential entrants that take draws from the productivity distribution) in step 1. Replace $\pi_i' = \pi_i$ with the free entry condition in each country: The expected value of taking a draw from the productivity distribution is exactly equal to the cost of taking a draw.

3   What's Next?

A large literature has developed taking on different aspects of these kinds of models. Melitz and Ottaviano (2008) move away from CES preferences, generating nonconstant markups in a tractable way. Bernard et al. (2003) combine a Ricardian model with Bertrand pricing to think about the distribution of firm productivity. [To do: Add more.]

Axtell, Robert L. (2001). "Zipf Distribution of U.S. Firm Sizes." *Science* 293 (5536), pp. 1818–20.

Bernard, Andrew B. et al. (2003). "Plants and Productivity in International Trade." *American Economic Review* 93 (4), pp. 1268–1290.

Chaney, Thomas (2008). "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *The American Economic Review* 98 (4), pp. 1707–1721.

Eaton, Jonathan and Samuel S. Kortum (2002). "Technology, Geography, and Trade." *Econometrica* 70 (5), pp. 1741–1779.

Hopenhayn, Hugo A (1992). "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5), pp. 1127–50.

Krugman, Paul R. (1980). "Scale Economies, Product Differentiation, and the Pattern of Trade." *American Economic Review* 70 (5), pp. 950–59.

Melitz, Marc J. (2003). "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6), pp. 1695–1725.

Melitz, Marc J. and Gianmarco I. P. Ottaviano (2008). "Market Size, Trade, and Productivity." *Review of Economic Studies* 75 (1), pp. 295–316.

Roberts, Mark J. and James R. Tybout (1997). "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs." *American Economic Review* 87 (4), pp. 545–64.