American Economic Review 2015, 105(10): 3183–3221
http://dx.doi.org/10.1257/aer.20120549

Competition, Markups, and the Gains from International Trade[†]

By Chris Edmond, Virgiliu Midrigan, and Daniel Yi Xu*

We study the procompetitive gains from international trade in a quantitative model with endogenously variable markups. We find that trade can significantly reduce markup distortions if two conditions are satisfied: (*i*) *there is extensive misallocation*, *and* (*ii*) *opening to trade exposes hitherto dominant producers to greater competitive pressure. We measure the extent to which these two conditions are satisfied in Taiwanese producer-level data. Versions of our model consistent with the Taiwanese data predict that opening up to trade strongly increases competition and reduces markup distortions by up to one-half*, *thus significantly reducing productivity losses due to misallocation.* (*JEL* D43, F12, F14, L13, L60, O47)

Can international trade significantly reduce product market distortions? We study this question in a quantitative trade model with endogenously variable markups. In such a model, markup dispersion implies that resources are *misallocated* and that aggregate productivity is low. By exposing producers to greater competition, international trade may reduce markup dispersion thereby reducing misallocation and increasing aggregate productivity. Our goal is to use producer-level data to quantify these *procompetitive effects*.

We study these procompetitive effects in the model of Atkeson and Burstein (2008). In this model, any given sector has a small number of producers who engage in *oligopolistic competition*. The demand elasticity for any given producer is decreasing in its market share and hence its markup is increasing in its market share. By reducing the market shares of dominant producers, international trade can

* Edmond: Department of Economics, University of Melbourne, 111 Barry Street, Carlton, VIC 3010, Australia (e-mail: cedmond@unimelb.edu.au); Midrigan: Department of Economics, New York University, 19 W. 4th Street, 6th Floor, New York, NY 10012, and NBER (e-mail: virgiliu.midrigan@nyu.edu); Xu: Department of Economics, Duke University, 419 Chapel Drive, Box 90097, Durham, NC 27708, and NBER (e-mail: daniel.xu@duke.edu). We thank four anonymous referees for valuable comments and suggestions. We have also benefited from discussions with Fernando Alvarez, Costas Arkolakis, Andrew Atkeson, Ariel Burstein, Vasco Carvalho, Andrew Cassey, Arnaud Costinot, Jan De Loecker, Dave Donaldson, Ana Cecilia Fieler, Oleg Itskhoki, Phil McCalman, Markus Poschke, Andrés Rodríguez-Clare, Barbara Spencer, Iván Werning, and from numerous conference and seminar participants. We also thank Andres Blanco, Sonia Gilbukh, Jiwoon Kim, and Fernando Leibovici for their excellent research assistance. We thank the National Science Foundation for financial support under grant SES-1156168. Edmond also thanks the Australian Research Council for financial support under grant DP-150101857. Midrigan also thanks the Alfred P. Sloan foundation for financial support. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. Midrigan is a consultant at the Federal Reserve Bank of Minneapolis.

[†] Go to http://dx.doi.org/10.1257/aer.20120549 to visit the article page for additional materials and author disclosure statement(s).

reduce markups and markup dispersion. The Atkeson and Burstein (2008) model is particularly useful for us because it implies a *linear* relationship between (inverse) producer-level markups and market shares, which in turn makes the model straightforward to parameterize.

We find that trade can significantly reduce markup distortions if two conditions are satisfied: (i) there is extensive misallocation, and (ii) international trade does in fact put producers under greater competitive pressure. The first condition is obvious—if there is no misallocation, there is no misallocation to reduce. The second condition is more subtle. Trade has to increase the degree of effective competition prevailing amongst producers within a market. If both domestic and foreign producers have similar productivities within a given sector, then opening to trade exposes them to genuine *head-to-head competition* that reduces market power thereby reducing markups and markup dispersion. By contrast, if there are large cross-country differences in productivity within a given sector, then opening to trade may allow producers from one country to substantially increase their market share in the other country, thereby increasing markups and markup dispersion so that the procompetitive "gains" from trade are negative.

We quantify the model using 7-digit Taiwanese manufacturing data. We use this data to discipline two key determinants of the extent of misallocation: (i) the elasticity of substitution across sectors, and (ii) the equilibrium distribution of producer market shares. The elasticity of substitution across sectors plays a key role because it determines the extent to which producers facing little competition in their own sector can raise markups. We pin down this elasticity by requiring that our model fits the cross-sectional relationship between measures of markups and market shares that we observe in the Taiwanese data. We pin down the parameters of the producer-level productivity distribution and fixed costs of operating and exporting by requiring that the model reproduces key moments of the distribution of market shares within and across sectors in the Taiwanese data.

The Taiwanese data feature a large amount of both dispersion and concentration in producer-level market shares and a strong relationship between market shares and measured markups. Interpreted through the lens of the model, this implies a significant amount of misallocation and hence the possibility of significant productivity gains from reduced markup distortions.

Given this misallocation, the model predicts large procompetitive gains if, within a given sector, domestic producers and foreign producers have relatively similar levels of productivity so that more trade increases the degree of competition prevailing among producers. This feature of the model is largely determined by the *cross-country correlation* in sectoral productivity. We choose the amount of correlation in sectoral productivity so that the model reproduces standard estimates of the elasticity of trade flows with respect to changes in variable trade costs. As the amount of correlation increases, there is less cross-country variation in producers' productivity. Consequently, small changes in trade costs have relatively larger effects on trade flows—in short, the *trade elasticity* is increasing in the amount of cross-country correlation. To match standard estimates of the trade elasticity, the benchmark model requires a relatively high 0.94 cross-country correlation in sectoral draws. This high correlation also allows the model to reproduce the strong positive relationship between a sector's share of domestic sales and its share of imports that we observe in the data—i.e., reproduces the fact that sectors with relatively large, productive firms are also sectors with relatively large import shares.

Given this high degree of correlation, opening to trade indeed reduces markup dispersion and increases aggregate productivity. For the benchmark model, calibrated to Taiwan's import share, opening to trade reduces markup distortions by about one-fifth and increases aggregate productivity by 12.4 percent relative to autarky. In short we find that, yes, opening to trade can lead to a quantitatively significant reduction in misallocation. We also find that these procompetitive effects are strongest near autarky—the procompetitive effects are more important for an economy opening from autarky to a 10 percent import share than for an economy increasing its openness from a 10 to 20 percent import share.

In the model, a given producer's productivity has both a sector-specific component and an idiosyncratic component, both drawn from Pareto distributions. In our benchmark model, the sectoral draws are correlated across countries while the idiosyncratic draws are not. We consider an extension of the model in which the idiosyncratic draws are also correlated across producers in a given sector in different countries. This extension is motivated by the observation that sectors with high concentration amongst domestic producers are also sectors with high import penetration. While our benchmark model cannot reproduce this feature of the data, our extension with correlated idiosyncratic draws can. This extension predicts an even larger role for trade in reducing markup distortions because countries import more of exactly those goods for which the domestic market is more distorted. In this version of the model, trade eliminates about one-third of the productivity losses from misallocation.

We consider a number of robustness checks on our benchmark model—including allowing for heterogeneity in sector-level tariffs, introducing labor market distortions, and changing the mode of competition from Cournot to Bertrand, amongst others. Our main findings are robust to these alternative specifications. We also study an extension of the model in which we introduce capital and elastic labor supply and show that the procompetitive gains from trade are even larger. Finally, we study a version of the model with free entry and show that versions of the free-entry model that reproduce the salient features of the Taiwanese data continue to predict significant procompetitive gains from trade.

Markups, *Misallocation*, *and Trade*.—Recent papers by Restuccia and Rogerson (2008), Hsieh and Klenow (2009), and others show that misallocation of factors of production can substantially reduce aggregate productivity. We focus on the role of markup variation as a source of misallocation.[1] We find that, by reducing markup dispersion, trade can play a powerful role in reducing misallocation and can thereby increase aggregate productivity.

The possibility that opening an economy to trade may lead to welfare gains from increased competition is, of course, one of the oldest ideas in economics. But standard quantitative trade models, such as the perfect competition model of Eaton and

[1] Two closely related papers are Peters (2013), who considers endogenous markups, as we do, in a closed economy quality-ladder model of endogenous growth and Epifani and Gancia (2011) who consider an open economy model but with exogenous markup dispersion.

Kortum (2002) or the monopolistic competition models with constant markups of Melitz (2003) and Chaney (2008), cannot capture this procompetitive intuition.

Perhaps more surprisingly, existing trade models that *do* feature variable markups do not generally predict procompetitive gains. For example, the Bernard et al. (2003) model of Bertrand competition results in an endogenous distribution of markups, that, due to specific functional form assumptions, is *invariant* to changes in trade costs and has exactly *zero* procompetitive gains.[2] Similarly, in the monopolistic competition models with non-CES demand[3] studied by Arkolakis et al. (2012), the markup distribution is likewise invariant to changes in trade costs and there are in fact *negative* procompetitive "gains" from trade.

The reason models with variable markups yield conflicting predictions regarding the procompetitive gains from trade is that, as emphasized by Arkolakis et al. (2012), what really matters for these effects is the *joint distribution* of markups and employment. The response of this joint distribution to a reduction in trade costs depends on the parameterization of the model, and in particular the amount of cross-country correlation in productivity draws. We show that versions of our model with low correlation do indeed predict negative procompetitive gains. But such parameterizations also imply both (i) low aggregate trade elasticities, and (ii) a weak or negative relationship between a sector's share of domestic sales and its share of imports, and thus are inconsistent with empirical evidence.

Empirical Literature on Markups and Trade.—There is a large empirical literature on producer markups and trade. Important early examples include Levinsohn (1993); Harrison (1994); and Krishna and Mitra (1998). Tybout (2003) reviews this literature and concludes that "in *every* country studied, relatively high sector-wide exposure to foreign competition is associated with lower price-cost margins, and the effect is concentrated in *larger* plants." More recently, Feenstra and Weinstein (2010) infer large markup reductions from observed changes in US market shares from 1992–2005. De Loecker et al. (2014) study the effects of India's tariff reductions on both final goods and inputs and find that the *net* effect was in fact an *increase* in markups—because input tariffs fell, so did the costs of final goods producers. When they condition on the effects of trade liberalization through inputs, however, De Loecker et al. find that the markups of final goods producers fall. Their results are thus consistent with our benchmark model.

There are important conceptual differences between the effects of trade in this literature and procompetitive gains through *reduced misallocation*. Documenting changes in the domestic markup distribution following a trade liberalization does not tell us whether misallocation has gone down or not. Again, what matters for misallocation is the response of the joint distribution of employment and markups of all producers, including exporters.

[2] An important contribution by De Blas and Russ (2015) extends Bernard et al. (2003) to allow for a finite number of producers in a given sector so that, as in our model, the distribution of markups varies in response to changes in trade costs. Holmes, Hsu, and Lee (2014) study the impact of trade on productivity and misallocation in this setting. Relative to these theoretical papers, as well as to Devereux and Lee (2001) and Melitz and Ottaviano (2008), our main contribution is to *quantify* the procompetitive gains from trade using micro data.

[3] Special cases of which include the non-CES demand systems used by Krugman (1979); Feenstra (2003); Melitz and Ottaviano (2008); and Zhelobodko et al. (2012).

Trade Flows and the Gains from Trade.—Our focus on the gains from trade is related to the work of Arkolakis, Costinot, and Rodríguez-Clare (2012)—henceforth, ACR—who show that the total gains from trade are identical in a large class of models and are summarized by the aggregate trade elasticity. Interestingly, we find that for our benchmark parameterization the ACR formula in fact provides an excellent approximation to the total gains from trade in our setup with variable markups.

That said, while the total gains from trade in our benchmark parameterization are well-approximated by the ACR formula, our model nevertheless predicts important procompetitive gains from trade. That is, opening up to trade substantially reduces markup distortions. Moreover, our model predicts that the gains from trade for two otherwise identical countries are larger for a country that has not yet reformed its product markets. Such differences in product market distortions are endogenously reflected in differences in the aggregate trade elasticity itself—which is precisely why the ACR formula can capture these procompetitive effects. Put differently, there can be important procompetitive gains in our model even when the ACR formula works well.[4]

The remainder of the paper proceeds as follows. Section I presents the model. Section II gives an overview of the data and Section III explains how we use that data to quantify the model. Section IV presents our benchmark results on the gains from trade. Section V conducts a number of robustness checks. Section VI presents results for two extensions of our benchmark model, (i) trade between asymmetric countries, and (ii) free entry and an endogenous number of competitors per sector. Section VII concludes.

I. Model

The economy consists of two symmetric countries, Home and Foreign. In keeping with standard assumptions in the trade literature, we assume a static environment with a single factor of production, labor, that is in inelastic supply and immobile between countries. We focus on describing the Home country in detail. We indicate Foreign variables with an asterisk.

A. *Final Good Producers*

Perfectly competitive firms in each country produce a homogeneous final consumption good Y using inputs $y(s)$ from a continuum of *sectors*

$$Y = \left(\int_0^1 y(s)^{\frac{\theta-1}{\theta}}\, ds\right)^{\frac{\theta}{\theta-1}}, \tag{1}$$

where $\theta > 1$ is the elasticity of substitution *across* sectors $s \in [0, 1]$. Importantly, each sector consists of a *finite* number of domestic and foreign intermediate

[4] To be clear, we measure the procompetitive gains as the reduction in misallocation induced by opening to trade. This is different from Arkolakis et al. (2012) who measure the procompetitive gains as the difference between the gains from trade in a given model with variable markups and the gains predicted by the ACR formula. This explains why we find "positive procompetitive effects" in our benchmark model even when the total gains from trade are well-approximated by the ACR formula.

producers. In sector s, output is produced using $n(s) \in \mathbb{N}$ domestic and $n(s)$ imported intermediate inputs

$$y(s) = \left(\sum_{i=1}^{n(s)} y_i^H(s)^{\frac{\gamma-1}{\gamma}} + \sum_{i=1}^{n(s)} y_i^F(s)^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}}, \tag{2}$$

where $\gamma > \theta$ is the elasticity of substitution across goods i *within* a particular sector $s \in [0, 1]$.

In our benchmark model, the number of potential producers $n(s)$ in sector s is exogenous and the same in both countries. In Section VI below we consider an extension of the benchmark model with *free entry* that makes the number of producers endogenous and varying across countries.[5]

B. *Intermediate Goods Producers*

Intermediate good producer i in sector s produces output using labor

$$y_i(s) = a_i(s) l_i(s), \tag{3}$$

where producer-level productivity $a_i(s)$ is drawn from a distribution that we discuss in detail in Section III below.

Trade Costs.—An intermediate good producer sells output to final good producers located in both countries. Let $y_i^H(s)$ denote the amount sold by a Home intermediate good producer to Home final good producers and similarly let $y_i^{*H}(s)$ denote the amount sold by a Home intermediate good producer to Foreign final good producers. The resource constraint for Home intermediate good producers is

$$y_i(s) = y_i^H(s) + \tau y_i^{*H}(s), \tag{4}$$

where $\tau \geq 1$ is an iceberg trade cost, i.e., $\tau y_i^{*H}(s)$ must be shipped for $y_i^{*H}(s)$ to arrive abroad. Foreign intermediate producers face symmetric trade costs. We let $y_i^*(s)$ denote their output and note that the resource constraint facing Foreign intermediate producers is

$$y_i^*(s) = \tau y_i^F(s) + y_i^{*F}(s), \tag{5}$$

where $y_i^{*F}(s)$ denotes the amount sold by a Foreign intermediate good producer to Foreign final good producers and $y_i^F(s)$ denotes the amount sold by a Foreign intermediate good producer to Home final good producers.

Demand for Intermediate Inputs.—Final good producers buy intermediate goods from Home producers at prices $p_i^H(s)$ and from Foreign producers at prices $p_i^F(s)$.

[5] In the online Appendix we also report results for a version of our model where the number of potential Home and Foreign producers per sector remain exogenous but are uncorrelated across countries.

Consumers buy the final good at price P. A final good producer chooses intermediate inputs $y_i^H(s)$ and $y_i^F(s)$ to maximize profits,

$$PY - \int_0^1 \left(\sum_{i=1}^{n(s)} p_i^H(s) y_i^H(s) + \tau \sum_{i=1}^{n(s)} p_i^F(s) y_i^F(s) \right) ds, \tag{6}$$

subject to (1) and (2). The solution to this problem gives the demand functions:

$$y_i^H(s) = \left(\frac{p_i^H(s)}{p(s)} \right)^{-\gamma} \left(\frac{p(s)}{P} \right)^{-\theta} Y, \tag{7}$$

and

$$y_i^F(s) = \left(\frac{\tau p_i^F(s)}{p(s)} \right)^{-\gamma} \left(\frac{p(s)}{P} \right)^{-\theta} Y, \tag{8}$$

where the aggregate and sectoral price indexes are

$$P = \left(\int_0^1 p(s)^{1-\theta} \, ds \right)^{\frac{1}{1-\theta}}, \tag{9}$$

and

$$p(s) = \left(\sum_{i=1}^{n(s)} \phi_i^H(s) p_i^H(s)^{1-\gamma} + \tau^{1-\gamma} \sum_{i=1}^{n(s)} \phi_i^F(s) p_i^F(s)^{1-\gamma} \right)^{\frac{1}{1-\gamma}}, \tag{10}$$

and where $\phi_i^H(s) \in \{0, 1\}$ is an indicator function that equals one if a producer operates in the Home market (its domestic market) and likewise $\phi_i^F(s) \in \{0, 1\}$ is an indicator function that equals one if a Foreign producer operates in the Home market (its export market).

Market Structure.—An intermediate good producer faces the demand system given by equations (7)–(10) and engages in *Cournot competition* within its sector.[6] That is, each individual firm chooses a quantity $y_i^H(s)$ or $y_i^{*H}(s)$ taking as given the quantity decisions of its competitors in sector s. Due to constant returns, the problem of a firm in its domestic market and its export market can be considered separately.

Fixed Costs.—There are fixed costs f_d and f_x of operating in the domestic and foreign market respectively. Both of these are denominated in units of domestic labor. A firm can choose to produce zero units of output for the domestic market to avoid paying the fixed cost f_d. Similarly, a firm can choose to produce zero units of output for the export market to avoid paying the fixed cost f_x. We introduce fixed operating costs in order to allow the model to match the lower tail of the distribution of firm size in the data.

[6] In Section V we solve our model under *Bertrand competition* and find similar results.

Domestic Market.—Taking the wage W as given, the problem of a Home firm in its domestic market can be written

$$(11) \qquad \pi_i^H(s) := \max_{y_i^H(s),\, \phi_i^H(s)} \left[\left(p_i^H(s) - \frac{W}{a_i(s)}\right) y_i^H(s) - W f_d\right] \phi_i^H(s),$$

subject to the demand system above. The solution to this problem is characterized by a price that is a markup over marginal cost

$$(12) \qquad p_i^H(s) = \frac{\varepsilon_i^H(s)}{\varepsilon_i^H(s) - 1} \frac{W}{a_i(s)},$$

where $\varepsilon_i^H(s) > 1$ is the demand elasticity facing the firm in its domestic market. With the nested CES demand system above and Cournot competition, it can be shown that this demand elasticity is a weighted *harmonic* average of the underlying elasticities of substitution θ and γ, specifically

$$(13) \qquad \varepsilon_i^H(s) = \left(\omega_i^H(s)\,\frac{1}{\theta} + (1 - \omega_i^H(s))\,\frac{1}{\gamma}\right)^{-1},$$

where $\omega_i^H(s) \in [0, 1]$ is the firm's share of sectoral revenue in its domestic market

$$(14) \qquad \omega_i^H(s) := \frac{p_i^H(s) y_i^H(s)}{\sum_{i=1}^{n(s)} p_i^H(s) y_i^H(s) + \tau \sum_{i=1}^{n(s)} p_i^F(s) y_i^F(s)} = \left(\frac{p_i^H(s)}{p(s)}\right)^{1-\gamma}.$$

For short, we refer to $\omega_i^H(s)$ as a Home firm's domestic *market share*.

Export Market.—The problem of a Home firm in its export market is essentially identical except that to export (operate abroad) it pays a fixed cost f_x rather than f_d so that its problem is

$$(15) \qquad \pi_i^{*H}(s) := \max_{y_i^{*H}(s),\, \phi_i^{*H}(s)} \left[\left(p_i^{*H}(s) - \frac{W}{a_i(s)}\right) y_i^{*H}(s) - W f_x\right] \phi_i^{*H}(s),$$

subject to the demand system abroad. Prices are again a markup over marginal cost

$$(16) \qquad p_i^{*H}(s) = \frac{\varepsilon_i^{*H}(s)}{\varepsilon_i^{*H}(s) - 1} \frac{W}{a_i(s)},$$

where $\varepsilon_i^{*H}(s) > 1$ is the demand elasticity facing the firm in its export market

$$(17) \qquad \varepsilon_i^{*H}(s) = \left(\omega_i^{*H}(s)\frac{1}{\theta} + \left(1 - \omega_i^{*H}(s)\right)\frac{1}{\gamma}\right)^{-1},$$

and where $\omega_i^{*H}(s) \in [0, 1]$ is the firm's share of sectoral revenue in its export market

$$(18) \qquad \omega_i^{*H}(s) := \frac{\tau p_i^{*H}(s) y_i^{*H}(s)}{\tau \sum_{i=1}^{n(s)} p_i^{*H}(s) y_i^{*H}(s) + \sum_{i=1}^{n(s)} p_i^{*F}(s) y_i^{*F}(s)}.$$

For short, we refer to $\omega_i^{*H}(s)$ as a Home firm's export market share.

Market Shares and Demand Elasticity.—In general, each firm faces a different, endogenously determined, demand elasticity. The demand elasticity is given by a weighted average of the within-sector elasticity γ and the across-sector elasticity $\theta < \gamma$. Firms with a small market share within a sector (within a given country) compete mostly with other firms in their own sector and so face a relatively high demand elasticity, closer to the within-sector γ. Firms with a large market share face relatively more competition from firms in other sectors than they do from firms in their own sector and so face a relatively low demand elasticity, closer to the across-sector θ. The markup a firm charges is an increasing *convex* function of its market share. An infinitesimal firm charges a markup of $\gamma/(\gamma - 1)$, the smallest possible in this model. At the other extreme, a pure monopolist charges a markup of $\theta/(\theta - 1)$, the largest possible in this model. Because of the convexity, a mean-preserving spread in market shares will increase the average markup.

The extent of *markup dispersion* across firms depends both on the *gap* between θ and γ and on the extent of dispersion in market shares. In the special case where $\theta = \gamma$, the demand elasticity is constant and independent of the dispersion in market shares and the model collapses to a standard trade model with constant markups. But if θ is substantially smaller than γ, then even a modest change in market share dispersion can have a large effect on markup dispersion and hence a large effect on aggregate productivity.

Notice also that a firm operating in both countries will generally have different market shares in each country and consequently face different demand elasticities and charge different markups in each country.

Market Shares and Markups.—The formula (13) for a firm's demand elasticity implies a linear relationship between a firm's *inverse markup* and its market share

$$(19) \qquad \frac{1}{\mu_i^H(s)} = \frac{\gamma - 1}{\gamma} - \left(\frac{1}{\theta} - \frac{1}{\gamma}\right) \omega_i^H(s),$$

where $\mu_i^H(s) := \varepsilon_i^H(s) / \big(\varepsilon_i^H(s) - 1\big)$ denotes the firm's gross markup from (12). Since $\theta < \gamma$, the coefficient on the market share $\omega_i^H(s)$ is negative. Within a sector s, firms with relatively high market shares have low demand elasticity and high markups. As discussed in Section III below, the strength of this relationship plays a key role in identifying plausible magnitudes for the gap between the elasticity parameters θ and γ.

Operating Decisions.—Each firm must pay a fixed cost f_d to operate in its domestic market and a fixed cost f_x to operate in its export market. A Home firm operates in its domestic market so long as

$$\left(p_i^H(s) - \frac{W}{a_i(s)}\right) y_i^H(s) \geq W f_d. \tag{20}$$

Similarly, a Home firm operates in its export market so long as

$$\left(p_i^{*H}(s) - \frac{W}{a_i(s)}\right) y_i^{*H}(s) \geq W f_x. \tag{21}$$

There are multiple equilibria in any given sector. Different combinations of firms may choose to operate, given that the others do not. As in Atkeson and Burstein (2008), within each sector s we place firms in the order of their physical productivity $a_i(s)$ and focus on equilibria in which firms sequentially decide on whether to operate or not: the most productive decides first (given that no other firm operates), the second most productive decides second (given that no other less productive firm operates), and so on.[7]

C. *Market Clearing*

In each country there is a representative consumer that inelastically supplies one unit of labor and consumes the final good. Let $l_i^H(s)$ denote the labor a Home firm uses in production for its domestic market and similarly let $l_i^{*H}(s)$ denote the labor a Home firm uses in production for its export market. The labor market clearing condition is then

$$\int_0^1 \left(\sum_{i=1}^{n(s)} \left(l_i^H(s) + f_d\right)\phi_i^H(s) + \sum_{i=1}^{n(s)} \left(l_i^{*H}(s) + f_x\right)\phi_i^{*H}(s)\right) ds = 1, \tag{22}$$

and the market clearing condition for the final good is simply $C = Y$.

D. *Aggregate Productivity and Markups*

Aggregation.—The quantity of final output in each country can be written

$$Y = A\tilde{L}, \tag{23}$$

[7] The exact ordering makes little difference quantitatively when we calibrate the model to match the strong concentration in the data. Productive firms always operate and unproductive ones never do, so the equilibrium multiplicity only affects the operating decisions of marginal firms that have a negligible effect on aggregates. Moreover, as we show in Section V below, our model's implications for markup dispersion are essentially unchanged when we set $f_d = f_x = 0$ so that all firms operate and the equilibrium is unique.

where A is the endogenous level of aggregate productivity and $\tilde{L}$ is the aggregate amount of labor employed *net of fixed costs*. Using the firms' optimality conditions and the market clearing condition for labor, it is straightforward to show that aggregate productivity is a *quantity-weighted* harmonic mean of firm productivities

$$(24)\qquad A = \left(\int_0^1 \left(\sum_{i=1}^{n(s)} \frac{1}{a_i(s)} \frac{y_i^H(s)}{Y} + \tau \sum_{i=1}^{n(s)} \frac{1}{a_i(s)} \frac{y_i^{*H}(s)}{Y}\right) ds\right)^{-1}.$$

Now denote the aggregate (economy-wide) markup by

$$(25)\qquad \mu := \frac{P}{W/A},$$

that is, aggregate price divided by aggregate marginal cost. It is straightforward to show that the aggregate markup is a *revenue-weighted* harmonic mean of firm markups

$$(26)\qquad \mu = \left(\int_0^1 \left(\sum_{i=1}^{n(s)} \frac{1}{\mu_i^H(s)} \frac{p_i^H(s) y_i^H(s)}{PY} + \tau \sum_{i=1}^{n(s)} \frac{1}{\mu_i^{*H}(s)} \frac{p_i^{*H}(s) y_i^{*H}(s)}{PY}\right) ds\right)^{-1},$$

where $\mu_i^H(s)$ denotes a Home firm's markup in its domestic market and $\mu_i^{*H}(s)$ denotes its markup in its export market (implied by equations (12) and (16), respectively).

Misallocation and Markup Dispersion.—In this model, dispersion in markups reduces aggregate productivity, as in the work of Restuccia and Rogerson (2008) and Hsieh and Klenow (2009). To understand this effect, first notice that the expression (24) for aggregate productivity can be written

$$(27)\qquad A = \left(\int_0^1 \left(\frac{\mu(s)}{\mu}\right)^{-\theta} a(s)^{\theta-1}\, ds\right)^{\frac{1}{\theta-1}},$$

where $\mu(s) := p(s)/(W/a(s))$ denotes the sector-level markup and where sector-level productivity is given by

(28)

$$a(s) = \left(\sum_{i=1}^{n(s)} \left(\frac{\mu_i^H(s)}{\mu(s)}\right)^{-\gamma} a_i(s)^{\gamma-1} \phi_i^H(s) + \tau^{1-\gamma} \sum_{i=1}^{n(s)} \left(\frac{\mu_i^F(s)}{\mu(s)}\right)^{-\gamma} a_i^*(s)^{\gamma-1} \phi_i^F(s)\right)^{\frac{1}{\gamma-1}}.$$

By contrast, the first-best level of aggregate productivity (the best attainable by a planner, subject to the trade cost τ) associated with an efficient allocation of resources is

$$A_{efficient} = \left(\int_0^1 a(s)^{\theta-1}\, ds\right)^{\frac{1}{\theta-1}}, \tag{29}$$

where sector-level productivity is

$$a(s) = \left(\sum_{i=1}^{n(s)} a_i(s)^{\gamma-1}\phi_i^H(s) + \tau^{1-\gamma}\sum_{i=1}^{n(s)} a_i^*(s)^{\gamma-1}\phi_i^F(s)\right)^{\frac{1}{\gamma-1}}, \tag{30}$$

with operating decisions $\phi_i^H(s)$, $\phi_i^F(s) \in \{0, 1\}$ as dictated by the solution to the planning problem. If there is no markup dispersion (as occurs, for example, if $\theta = \gamma$), then aggregate productivity from (27)–(28) is at its first-best level. Markup dispersion lowers aggregate productivity relative to the first-best because it induces an inefficient allocation of resources across producers (relative prices are not aligned with relative marginal costs).

E. *Trade Elasticity*

As emphasized by ACR, in standard trade models the gains from trade are largely determined by the elasticity of trade flows with respect to changes in trade costs. We follow standard practice in the trade literature and define this *trade elasticity* as

$$\sigma := \frac{d\log\frac{1-\lambda}{\lambda}}{d\log\tau}, \tag{31}$$

where λ denotes the aggregate share of spending on domestic goods,

$$\lambda := \frac{\int_0^1 \sum_{i=1}^{n(s)} p_i^H(s)y_i^H(s)\, ds}{\int_0^1 \left(\sum_{i=1}^{n(s)} p_i^H(s)y_i^H(s) + \tau\sum_{i=1}^{n(s)} p_i^F(s)y_i^F(s)\right) ds} = \int_0^1 \lambda(s)\omega(s)\, ds, \tag{32}$$

and where $\lambda(s)$ denotes the sector-level share of spending on domestically produced goods and $\omega(s) := (p(s)/P)^{1-\theta}$ is that sector's share of aggregate spending.

To derive an expression for the trade elasticity σ in our model, we begin with a simpler calculation, showing how trade flows respond to changes in *international relative prices*. In a standard model with constant markups, this would also give us the trade elasticity. But with variable markups it does not. With variable markups there is *incomplete pass-through*: a 1 percent fall in trade costs reduces the relative price of foreign goods by less than 1 percent. We then show how this simpler calculation needs to be modified to account for incomplete pass-through.

Response of Trade Flows to International Relative Prices.—Suppose all foreign prices uniformly change by a factor q (this may be because of changes in trade costs, productivity, or labor supply, etc.). With a bit of algebra it can be shown that, in our model, the elasticity of trade flows with respect to international relative prices is given by a weighted average of the underlying elasticities of substitution γ and θ, specifically[8]

$$\frac{d\log\frac{1-\lambda}{\lambda}}{d\log q} = \gamma\left(\int_0^1 \frac{\lambda(s)}{\lambda}\left(\frac{1-\lambda(s)}{1-\lambda}\right)\omega(s)\,ds\right) + \theta\left(1-\int_0^1 \frac{\lambda(s)}{\lambda}\left(\frac{1-\lambda(s)}{1-\lambda}\right)\omega(s)\,ds\right) - 1$$

so that

$$\frac{d\log\frac{1-\lambda}{\lambda}}{d\log q} = (\gamma-1)-(\gamma-\theta)\frac{\text{var}[\lambda(s)]}{\lambda(1-\lambda)}, \tag{33}$$

where $\text{var}[\lambda(s)]$ is the variance across sectors of the share of spending on domestic goods and λ is the aggregate share, as defined in (32). We refer to the term $\text{var}[\lambda(s)]/\lambda(1-\lambda)$ as our *index of import share dispersion*. Notice that this elasticity is generally less than $\gamma - 1$ and is decreasing in the index of import share dispersion. If there is no import share dispersion, $\lambda(s) = \lambda$ for all s, then $\text{var}[\lambda(s)] = 0$ and the elasticity is relatively high, equal to $\gamma - 1$. Intuitively, if all sectors have identical import shares then there is no across-sector reallocation of expenditure and a uniform reduction in the relative price of foreign goods symmetrically increases import shares within each sector, an effect governed by γ. At the other extreme, if import shares are binary, $\lambda(s) \in \{0, 1\}$, then $\text{var}[\lambda(s)] = \lambda(1-\lambda)$ and the elasticity is relatively low, equal to $\theta - 1$. Here there is *only* across-sector reallocation of expenditure and a uniform reduction in the relative price of foreign goods induces reallocation toward sectors with high import shares, an effect governed by θ.

In a standard model, with constant markups, we would have $d\log q = d\log\tau$ and so the formula for the elasticity of trade flows with respect to international relative prices in (33) would also give us the trade elasticity σ. But in our model, with variable markups, there is incomplete pass-through from changes in trade costs to changes in relative prices and we need to modify (33) to account for these effects.

[8] Our goal here is to obtain analytic results that aid in building intuition. To that end, in the following expressions we abstract from the extensive margin and hold the set of producers in each country fixed. We relax this assumption and determine the set of operating firms endogenously when we compute the trade elasticity σ in our model. It turns out that treating the set of producers as fixed is, quantitatively, a good approximation in our model. In particular, as we show in Section V below, the quantitative implications of our model are almost identical when there are no fixed costs and all producers operate in both countries.

Accounting for Incomplete Pass-Through.—To account for incomplete pass-through, begin by noting that at the sector level the responsiveness of trade flows to trade costs is

$$\frac{d\log\frac{1-\lambda(s)}{\lambda(s)}}{d\log\tau} = (\gamma-1)(1+\epsilon(s)),$$

where

$$\epsilon(s) := \sum_{i=1}^{n(s)} \frac{p_i^F(s)y_i^F(s)}{p^F(s)y^F(s)}\left(\frac{d\log\mu_i^F(s)}{d\log\tau}\right) - \sum_{i=1}^{n(s)} \frac{p_i^H(s)y_i^H(s)}{p^H(s)y^H(s)}\left(\frac{d\log\mu_i^H(s)}{d\log\tau}\right),$$

denotes the elasticity with respect to trade costs of Foreign markups relative to Home markups and where $p^F(s)y^F(s)$ and $p^H(s)y^H(s)$ denote spending on Foreign goods and spending on Home goods in sector s. In general, the relative markup elasticity $\epsilon(s)$ is *negative*—i.e., a reduction in trade costs tends to increase Foreign markups as their producers gain market share and to decrease Home markups as their producers lose market share.

The aggregate trade elasticity σ can then be written

$$\sigma = (\gamma-\theta)\left(\int_0^1 \frac{\lambda(s)}{\lambda}\left(\frac{1-\lambda(s)}{1-\lambda}\right)(1+\epsilon(s))\omega(s)\,ds\right) + (\theta-1)\left(\int_0^1 \left(\frac{1-\lambda(s)}{1-\lambda}\right)(1+\epsilon(s))\omega(s)\,ds\right). \tag{34}$$

Further Intuition.—To see how this relates to our simple expression in (33) above, notice that in the special case where the relative markup elasticity is the same in each sector, $\epsilon(s) = \epsilon$ for all s, equation (34) reduces to

$$\sigma = \left((\gamma-1) - (\gamma-\theta)\frac{\mathrm{var}[\lambda(s)]}{\lambda(1-\lambda)}\right)(1+\epsilon).$$

Comparing this with (33) we see that, for this special case, the trade elasticity σ is proportional to the elasticity with respect to international relative prices. In the further special case of $\gamma = \theta$, so that markups are constant, then $\epsilon = 0$ (there is complete pass-through) and the trade elasticity indeed coincides with the elasticity of trade flows with respect to international relative prices—in this case, both elasticities equal $\gamma - 1$. With variable markups, the trade elasticity is generally less than $\gamma - 1$, both because the elasticity with respect to international relative prices is less than $\gamma - 1$ and because the elasticity with respect to trade costs is less than that with respect to relative prices.

II. Data

We now describe the data we use. First we give a brief description of the Taiwanese dataset. We then highlight facts about producer concentration in this data that are crucial for our model's quantitative implications.

A. *Dataset*

We use the Taiwan Annual Survey of Manufacturing (Ministry of Economic Affairs, Taiwan 2000–2004) that reports data for the universe of establishments[9] engaged in production activities. Our sample covers the years 2000 and 2002–2004. The year 2001 is missing because in that year a separate census was conducted.

Product Classification.—The dataset we use has two components. First, an *establishment*-level component collects detailed information on operations, such as employment, expenditure on labor, materials and energy, and total revenue. Second, a *product*-level component reports information on revenues for each of the products produced at a given establishment. Each product is categorized into a 7-digit Standard Industrial Classification created by the Taiwanese Statistical Bureau. This classification at seven digits is comparable to the detailed 5-digit SIC product definition collected for US manufacturing establishments as described by Bernard, Redding, and Schott (2010). Panel A of online Appendix Table A1 gives an example of this classification, while panel B reports the distribution of 7-digit sectors within 4- and 2-digit industries. Most of the products are concentrated in the Chemical Materials, Industrial Machinery, Computer/Electronics, and Electrical Machinery industries.

Import Shares.—We supplement the survey with detailed import data at the harmonized HS-6 product level. We obtain the import data from the World Trade Organization and then match HS-6 codes with the 7-digit product codes used in the Annual Manufacturing Survey. This match gives us disaggregated import penetration ratios for each product category.

B. *Concentration Facts*

The amount of producer concentration in the Taiwanese manufacturing data is crucial for our model's quantitative implications.

Strong Concentration within Sectors.—We measure a producer's market share by their share of domestic sales revenue within a given 7-digit sector. Panel A of Table 1 shows that producers within a sector are highly concentrated. The top producer has a market share of around 40 to 45 percent.[10] The median inverse Herfhindhal (HH) measure of concentration is about 3.9, much lower than 10 or so producers that

[9] In the Taiwanese data, almost all firms are single-establishment. In the online Appendix we show that using firm-level data rather than establishment-level data makes almost no difference to our results. If anything, using establishments rather than firms *understates* the extent of concentration among producers, a key feature that determines the gains from trade in our model.

[10] We weight each sector by the sector's share of aggregate sales.

Table 1—Parameterization

	Data	Model		Data	Model
Panel A. Moments					
Within-sector concentration, domestic sales			Size distribution sectors, domestic sales		
Mean inverse hh	7.25	4.30	Fraction sales by top 0.01 sectors	0.26	0.21
Median inverse hh	3.92	3.79	Fraction sales by top 0.05 sectors	0.52	0.32
Mean top share	0.45	0.46	Fraction wages (same) top 0.01 sectors	0.11	0.22
Median top share	0.40	0.41	Fraction wages (same) top 0.05 sectors	0.32	0.33
Distribution of sectoral shares, domestic sales			Size distribution producers, domestic sales		
Mean share	0.04	0.05	Fraction sales by top 0.01 producers	0.41	0.33
Median share	0.005	0.005	Fraction sales by top 0.05 producers	0.65	0.61
p75 share	0.02	0.03	Fraction wages (same) top 0.01 producers	0.24	0.31
p95 share	0.19	0.27	Fraction wages (same) top 0.05 producers	0.47	0.57
p99 share	0.59	0.59			
SD share	0.11	0.12			
Across-sector concentration			Coefficients in regression of labor share on market share		
p10 inverse HH	1.17	1.70	Ratio of coefficients b_1/b_0	−0.78	−0.78
p50 inverse HH	3.73	3.79			
p90 inverse HH	13.82	7.66	Import and export statistics		
p10 top share	0.16	0.24	Aggregate fraction exporters	0.25	0.25
p50 top share	0.41	0.41	Aggregate import share	0.38	0.38
p90 top share	0.92	0.75	Trade elasticity	4.00	4.00
p10 number producers	2	3	Coefficient, share imports on share sales	0.81	0.55
p50 number producers	10	16	Index import share dispersion	0.38	0.26
p90 number producers	52	47	Index intraindustry trade	0.37	0.45

Panel B. Parameter values

γ	10.50	Within-sector elasticity of substitution
θ	1.24	Across-sector elasticity of substitution
ξ_x	4.58	Pareto shape parameter, idiosyncratic productivity
ξ_z	0.51	Pareto shape parameter, sector productivity
ζ	0.043	Geometric parameter, number producers per sector
f_d	0.004	Fixed cost of domestic operations
f_x	0.203	Fixed cost of export operations
τ	1.129	Gross trade cost
ρ	0.94	Kendall correlation for Gumbel copula

operate in a typical sector. The distribution of market shares is skewed to the right and extremely *fat-tailed*. The median market share of a producer is just 0.5 percent while the average market share is 4 percent. The ninety-fifth percentile accounts for only 19 percent of sales while the ninety-ninth percentile accounts for 59 percent of sales. The overall pattern that emerges is consistently one of very strong concentration. Although quite a few producers operate in any given sector, most of these producers are small and a few large producers account for the bulk of the sector's domestic sales.

Strong Unconditional Concentration.—Panel A of Table 1 also reports statistics on the distribution of sales revenue and the wage bill across sectors and across all producers. The top 1 percent of sectors alone accounts for 26 percent of aggregate sales and 11 percent of the aggregate wage bill. The top 5 percent of sectors accounts for about one-half of all sales and about one-third of the wage bill. This pattern is reproduced at the producer level. The top 1 percent of producers accounts for 41 percent of sales and 24 percent of the wage bill, the top 5 percent of producers accounts for nearly two-thirds of all sales and nearly one-half of the wage bill. Again, the overall pattern is thus of strong concentration both within and across sectors.

III. Quantifying the Model

In the model, the size of the gains from trade largely depends on two factors: (i) the extent of misallocation, and (ii) the responsiveness of that misallocation to changes in trade costs. In turn, these factors are largely determined by the joint distribution of productivity, both within and across countries, and on the elasticity of substitution parameters θ and γ. We discipline our model along these dimensions by requiring that it reproduces a number of stylized features of the data: the amount of concentration within and across sectors, the relationship between a producer's labor share and market share, standard estimates of the trade elasticity, and the amount of intraindustry trade.

We next discuss our choice of functional forms for the productivity distributions and the parameter values we use in our quantitative work. In our discussion, we build intuition for our identification strategy by discussing heuristically how each parameter of the model is pinned down by key features of the data. But to be clear, formally, our calibration procedure involves choosing simultaneously a vector of parameters that minimizes the weighted distance between a vector of model moments and their data counterparts.

A. *Parameterization*

Within-Country Productivity Distribution.—We assume that across sectors the number of producers $n(s) \in \mathbb{N}$ is drawn i.i.d. from a geometric distribution with parameter $\zeta \in (0, 1)$ so that $\text{Prob}[n] = (1 - \zeta)^{n-1}\zeta$ and the average number of producers per sector is $1/\zeta$. We assume that an individual producer's productivity $a_i(s)$ is the product of a sector-specific component and an idiosyncratic component

$$a_i(s) = z(s)x_i(s). \tag{35}$$

We assume $z(s) \geq 1$ is independent of $n(s)$ and is drawn i.i.d. from a Pareto distribution with shape parameter $\xi_z > 0$ across sectors. Within sector s, the $n(s)$ draws of the idiosyncratic component $x_i(s) \geq 1$ are i.i.d. Pareto across producers with shape parameter $\xi_x > 0$.

Cross-Country Productivity Distribution.—We assume that cross-country correlation in productivity arises through correlation in sectoral productivities. In particular, let $F_Z(z)$ denote the Pareto distribution of sector-specific productivities within each country and let $H_Z(z, z^*)$ denote the cross-country joint distribution of these sector-specific productivities. We write this cross-country joint distribution as

$$H_Z(z, z^*) = \mathcal{C}(F_Z(z), F_Z(z^*)), \tag{36}$$

where the *copula* $\mathcal{C}$ is the joint distribution of a pair of uniform random variables u, u^* on $[0, 1]$. This formulation allows us to first specify the marginal distribution $F_Z(z)$ so as to match within-country productivity statistics and to then use the copula function to control the pattern of dependence between z and z^*.

Specifically, we assume that the marginal distributions are linked by a *Gumbel copula*, a widely used functional form that allows for dependence even in the right tails of the distribution,

$$(37)\quad \mathcal{C}(u, u^*) = \exp\left(-\left[(-\log u)^{\frac{1}{1-\rho}} + (-\log u^*)^{\frac{1}{1-\rho}}\right]^{1-\rho}\right), \quad 0 \leq \rho \leq 1 .$$

When working with heavy-tailed distributions, it is standard to summarize dependence using a robust correlation coefficient known as "Kendall's tau" (Nelsen 2006). For the copula above, this corresponds to the parameter ρ. If $\rho = 0$, then the copula reduces to $\mathcal{C}(u, u^*) = uu^*$ so that the draws are independent. If $\rho \rightarrow 1$ then the copula approaches $\mathcal{C}(u, u^*) = \min[u, u^*]$ so that the draws are perfectly dependent. Once the within-country distribution $F_Z(z)$ has been specified, the single parameter ρ pins down the joint distribution $H_Z(z, z^*)$.

Finally, let $F_X(x)$ denote the Pareto distribution of idiosyncratic productivities within each sector and let $H_X(x, x^*)$ denote the associated joint distribution. For our benchmark model we assume these are independent across countries so that $H_X(x, x^*) = F_X(x)F_X(x^*)$.

B. *Calibration*

We simultaneously choose a vector of nine parameters

$$\xi_x, \ \xi_z, \ \zeta, f_d, f_x, \ \tau, \ \rho, \ \gamma, \ \theta$$

to minimize the distance between a large number of model moments and their counterparts in the Taiwanese data. Panel A of Table 1 of reports the moments we target and the counterparts for our benchmark model. Panel B reports the parameter values that achieve this fit. We now provide intuition for how each parameter is determined by key features of the data.

Number of Producers, Productivity, and Fixed Cost of Operating.—The parameters ζ, ξ_z, ξ_x governing the within-country productivity distribution and the fixed cost f_d of operating in the domestic market are mainly determined by the pattern of concentration in the Taiwanese data. Intuitively, the geometric parameter ζ mainly determines the median number of producers per sector, the Pareto shape parameters ξ_z, ξ_x mainly determine the amount of concentration across-sectors and within-sectors, respectively, and the fixed cost f_d mainly influences the median size of producers.

Our model successfully reproduces the amount of concentration in the data. Within a given sector, the largest producer accounts for an average 46 percent of that sector's domestic sales (45 percent in the data). The model also reproduces the heavy concentration in the tails of the distribution of market shares with the ninety-ninth percentile share being 59 percent in both model and data. Moreover, the model also produces a fat-tailed size distribution of sectors and a fat-tailed size distribution of producers. The ninety-ninth percentile of sectors accounts for 21 percent of domestic sales (26 percent in the data) while the ninety-ninth percentile of producers accounts for 33 percent of domestic sales (41 percent in the data). The median number of producers per sector is a little too high (16 in the model, 10 in the data) but the model reproduces well the dispersion in the number of producers per sector (the tenth percentile is 3 producers in the model and 2 in the data, the ninetieth percentile is 47 producers in the model and 52 in the data).

The within-country joint distribution of productivity $a_i(s) = z(s)x_i(s)$ that generates this concentration is likewise very fat-tailed. This mostly comes from the sectoral productivity effect, $z(s)$, which has a Pareto shape parameter $\xi_z = 0.51$. By contrast, the idiosyncratic productivity effect, $x_i(s)$, has relatively thin tails with a Pareto shape parameter $\xi_x = 4.58$. The fixed cost of operating domestically is quite small, $f_d = 0.004$. This is about 0.26 percent of the average domestic producer's profits and 0.08 percent of their wage bill. This value of f_d is required for the model to reproduce the median size of producers we observe in the data. If f_d were smaller, the median sectoral share of producers would be much smaller than the 0.5 percent observed in the data.

Trade Costs.—The proportional trade cost τ and the fixed cost of operating in the export market are mainly pinned down by the requirement that the model reproduces Taiwan's aggregate import share of 0.38 and aggregate fraction of firms that export of 0.25. The model achieves this with a trade cost of $\tau = 1.129$ (i.e., 1.129 units of a good must be shipped for 1 unit to arrive) and a quite large fixed cost of operating in the export market, $f_x = 0.203$. This is about 3.3 percent of the average exporter's profits and 1.0 percent of their wage bill.

Cross-Country Correlation.—The copula parameter ρ governing the degree of cross-country correlation in sectoral productivity is mostly pinned down by the requirement that our model produces realistic values for (i) the cross-sectional relationship between sector import shares and sector domestic size, (ii) the amount of import share dispersion, and (iii) the amount of intraindustry trade. For all these statistics we simply target their counterparts in the Taiwanese manufacturing data.

The implied value of ρ is equal to 0.94 so that there is a high degree of correlation in productivity draws across countries. The model produces a somewhat small elasticity of sectoral import shares on a sector's overall sales share (0.55 in the model versus 0.81 in the data), suggesting that, if anything, we understate the amount of covariation between import shares and sector size. On the other hand, the model somewhat overstates the degree of intraindustry trade: it overpredicts the Grubel-Lloyd index (0.45 in the model versus 0.37 in the data), and produces too little import share dispersion (0.26 in the model versus 0.38 in the data). Since we use a single parameter, ρ, to jointly target these and other statistics, including the aggregate trade elasticity, our algorithm does not match any of these perfectly, but tries to match the entire set of moments as best as possible.

We discuss the sensitivity of our results to our choice of ρ below.

Elasticities of Substitution.—In our model, the elasticities θ and γ have important implications for both the extent to which changes in trade costs affect trade flows—the trade elasticity, as well as for the extent to which greater market power within a given industry allows a producer to increase its markups. Since both of these ingredients are important for our model's predictions about how trade changes allocations, we ask the model to match both of these features of the data.

In particular, we require that the model matches the consensus estimate of the aggregate trade elasticity of $\sigma = 4$ (see Simonovska and Waugh 2014), as well as the extent to which a producer's labor share (which in the model is related one-for-one to markups) changes with the producer's market share. In particular, we use an indirect inference approach and require that the model reproduces the *ratio* of the coefficients b_1 to b_0 in regressions of the form

$$\frac{Wl_i(s)}{p_i(s)y_i(s)} = b_0 + b_1\,\omega_i(s), \tag{38}$$

where the dependent variable is the producer's labor share and the independent variable is the producer's market share. To see why the ratio b_1/b_0 is informative about θ and γ, recall that in the model the labor share is inversely related to markups and moreover, the model implies that

$$\theta = \left(\frac{1}{\gamma} - \frac{b_1}{b_0}\left(\frac{\gamma - 1}{\gamma}\right)\right)^{-1} \tag{39}$$

so the ratio imposes a restriction on what value θ can take given a particular value of γ and vice-versa.[11]

In the Taiwanese data we obtain an intercept $b_0 = 0.64$ and slope $b_1 = -0.50$ so that we require our model to match the ratio $b_1/b_0 = -0.78$. Our calibration procedure chooses $\gamma = 10.5$ and hence $\theta = 1.24$ to satisfy this restriction as well as to match a trade elasticity of $\sigma = 4$. Intuitively, the data show a strong comovement between market shares and labor shares and the model can only match this if the two elasticities γ and θ are sufficiently far apart so that large producers face a low demand elasticity and are able to charge high markups. Moreover, since the pass-through of trade costs in prices is far below unity, the model requires a fairly high sectoral elasticity γ to match a trade elasticity of 4.

Notice finally that $\xi_z = 0.51 > \theta - 1 = 0.24$ at our calibrated parameter values, so that aggregate quantities are bounded despite the fact that the Pareto shape of sectoral productivities is less than 1.

Alternative Markup Estimates.—In our model, as is standard in the trade literature, labor is the only factor of production and a producer's inverse labor share *is* its markup. But in comparing our model's implications for markups to the data, it is important to recognize that, in general, factor shares differ across producers not only because of markup differences but also because of differences in the technology with which they operate. To control for this potential source of heterogeneity, in the online Appendix we follow De Loecker and Warzynski (2012) and

[11] Taking the model at face value, we could in fact simply back out the value of γ from the intercept b_0 as in (19) and then recover θ using b_1. But this approach is only valid if the production function is exactly linear in labor. Moreover, given how important the trade elasticity is for the model's aggregate implications, we follow the approach in the trade literature and choose parameters to match a given trade elasticity. See the online Appendix for more details.

TABLE 2—MARKUPS IN DATA AND MODEL

		Benchmark		Bertrand	
	Data	Taiwan	Autarky	Taiwan	Autarky
Panel A. Markup moments					
Aggregate markup		1.31	1.35	1.21	1.24
Unconditional markup distribution					
Mean	1.13	1.15	1.14	1.11	1.12
p50	1.11	1.11	1.11	1.11	1.11
p75	1.12	1.13	1.12	1.11	1.11
p90	1.14	1.21	1.18	1.12	1.12
p95	1.18	1.31	1.27	1.15	1.15
p99	1.41	1.68	1.64	1.33	1.37
SD log	0.06	0.08	0.09	0.04	0.08
log p95/p50	0.06	0.17	0.14	0.04	0.04
Across-sector markup distribution					
Mean	1.29	1.37	1.54	1.22	1.26
p50	1.18	1.30	1.31	1.18	1.18
p75	1.30	1.40	1.45	1.24	1.26
p90	1.50	1.61	1.83	1.39	1.49
p95	1.77	1.84	2.50	1.54	2.17
p99	2.76	2.22	5.25	2.09	5.25
SD log	0.18	0.14	0.31	0.12	0.32
log p95/p50	0.41	0.35	0.65	0.27	0.61
Panel B. Aggregate implications					
Import share	0.38	0.38	0	0.38	0
Fraction exporters	0.25	0.25	0	0.25	0
TFP loss, percent		7.0	9.0	2.1	4.6
Gains from trade, percent		12.4	—	13.8	—
Procompetitive gains, percent		2.0	—	2.5	—

Notes: The benchmark model features Cournot competition. TFP losses are the percentage gap between the level of aggregate productivity and the first-best level of aggregate productivity associated with the planning allocation (subject to the same trade costs). The gains from trade are the percentage change in aggregate productivity relative to autarky. The procompetitive gains from trade are the percentage change in aggregate productivity less the percentage change in first-best productivity.

use state-of-the-art IO methods to estimate markups that are purged of producer and sector-level differences in technology. Reassuringly, we find that our estimates of θ and γ are essentially unchanged if we use this alternative measure of markups. In particular, we find almost identical implications for the gains from international trade. See the online Appendix for a more detailed discussion of these alternative markup estimates and their implications.

C. *Markup Distribution*

Table 2 reports moments of the distribution of markups $\mu_i(s)$ in our benchmark model and their counterparts in the data (measured as the inverse of the fitted values of the labor share from (38)). We compare these to an economy that is identical except that we shut down international trade.

Panel A of Table 2 reports moments of the unconditional markup distribution, pooling over all sectors. The benchmark model implies an average markup of 1.15, a median markup of 1.11 (only slightly above $\gamma/(\gamma - 1) = 1.105$), and a standard deviation of log markups of 0.08. Moreover, as in the data larger producers have considerably higher markups. The ninety-fifth percentile markup is 1.31 (compared to 1.18 in the data) and the ninety-ninth percentile markup is 1.68 (compared to 1.41 in the data)—though note that these are still short of the $\theta/(\theta - 1) = 5.25$ markup a pure monopolist would charge in our model. Because large producers charge higher markups, the aggregate markup, which is a revenue-weighted harmonic average of the individual markups, is 1.31—much higher than the simple average.

Let $\mu(s) = p(s)/(W/a(s))$ denote the aggregate markup in sector s. This sector-level markup $\mu(s)$ is likewise a revenue-weighted harmonic average of the producer-level markups $\mu_i(s)$ within that sector. Both in the model and in the data, these sector-level markups $\mu(s)$ are larger and more dispersed than their producer-level counterparts $\mu_i(s)$. In the model, the median sectoral markup is 1.30 as opposed to 1.11 for producers while the ninety-ninth percentile sectoral markup is 2.22 as opposed to 1.68 for producers. In short, markup dispersion across sectors is at least as great as markup dispersion within sectors. Note also that the model fails to replicate the full extent of the across-sector variation in markups, especially in the tails. The ninety-ninth percentile markup in the data is 2.76, as opposed to 2.22 in the model. Since the actual dispersion in markups across sectors is larger than in the model, this suggests that we are, if anything, *understating* the true losses from markup dispersion.

Now consider what happens when we shut down all international trade, which we report in the column labeled Autarky. The *unconditional* markup distribution hardly changes. The median markup is unchanged and, if anything, there is a slight increase in the unconditional markup dispersion. Nonetheless, there is substantially more misallocation under autarky. As shown in panel B of Table 2, the benchmark economy implies aggregate productivity 7 percent below the first-best level of productivity associated with the planning allocation. Under autarky, the economy is 9 percent below the first-best. Hence the extent of misallocation is considerably worse under autarky.

As emphasized by Arkolakis et al. (2012), moments of the unconditional markup distribution are a poor guide to evaluating the procompetitive gains from trade—as they show, in several important theoretical benchmarks, the unconditional markup distribution is *invariant* to the level of trade costs. Instead, what matters is the *joint distribution* of markups and employment across producers. In our benchmark model, opening to trade dramatically reduces the markups of the largest producers where most employment is concentrated.[12] This can be seen by comparing the moments of the sectoral markup distribution to its counterpart under autarky. Under autarky, the ninety-ninth percentile sectoral markup is 5.25—i.e., these sectors are pure monopolies—but with trade, the ninety-ninth percentile markup falls to 2.22 as these monopolists lose substantial market share to foreign competition. The standard deviation of log sectoral markups falls by about one-half, from 0.31 to 0.14, with much of this reduction coming from a fall in the markups of dominant producers that account for a large share of employment. As a result of this, misallocation falls from 9 percent to 7 percent.

[12] The response of the joint distribution of markups and employment to a change in trade costs depends sensitively on details of the parameterization of the model. We discuss this at length below.

TABLE 3—GAINS FROM TRADE

Change in import share	0 to 10	10 to 20	20 to 30	30 to Taiwan	0 to Taiwan
Panel A. Benchmark model					
Change TFP, percent	3.4	2.7	3.3	3.0	12.4
Change first-best TFP, percent	1.8	2.5	3.2	3.0	10.4
Procompetitive gains, percent	1.7	0.3	0.1	0.0	2.0
Misallocation relative to autarky	0.81	0.79	0.78	0.78	0.78
Change aggregate markup, percent	−1.9	−0.6	−0.4	−0.1	−2.9
Domestic	−1.6	−0.6	−0.4	−0.3	−2.9
Import	16.6	−0.1	0.4	0.2	17.1
Change markup dispersion, percent	−1.7	−0.2	1.1	−0.1	−0.9
Domestic	−1.9	−0.4	1.0	−0.4	−1.7
Import	10.3	−0.1	0.0	0.1	10.3
Trade elasticity (ex post)	4.2	4.1	4.0	4.0	4.0
ACR gains, percent	2.5	2.9	3.3	2.8	11.7
Panel B. Alternative model with correlated $x_i(s)$, $x_i^*(s)$					
Change TFP, percent	3.3	2.6	3.1	3.0	12.0
Change first-best TFP, percent	1.6	2.2	2.9	2.9	9.6
Procompetitive gains, percent	1.7	0.3	0.2	0.1	2.4
Misallocation relative to autarky	0.81	0.77	0.75	0.73	0.73
Change aggregate markup, percent	−2.0	−0.8	−0.5	−0.2	−3.5
Domestic	−1.7	−0.6	−0.5	−0.4	−3.2
Import	15.7	−0.2	0.3	0.3	16.1
Change markup dispersion, percent	−0.4	−1.7	0.9	−1.2	−2.4
Domestic	−0.3	−2.0	0.9	−1.7	−3.1
Import	8.1	0.0	0.1	0.2	8.4
Trade elasticity (ex post)	4.2	4.2	4.1	4.0	4.0
ACR gains, percent	2.5	2.8	3.3	2.8	11.7

Notes: Panel A shows the gains from trade for our benchmark model. Panel B shows the gains from trade for our alternative model with correlation in idiosyncratic draws $x_i(s)$, $x_i^*(s)$ chosen to match the cross-sectional relationship between import penetration and domestic producer concentration, as discussed in the main text. For our benchmark model $x_i(s)$, $x_i^*(s)$ are independent and there is cross-country correlation in productivity only through correlation in sectoral productivity $z(s)$, $z^*(s)$. Markup dispersion measured by the standard deviation of log markups.

IV. Gains from Trade

We now calculate the aggregate productivity gains from trade in our benchmark model. As in ACR, we focus on the gains due to a permanent reduction in trade costs τ. We then ask our key question: to what extent does international trade reduce misallocation due to markups?

A. *Total Gains from Trade*

We measure the gains from trade by the percentage change in aggregate productivity from one equilibrium to another (the response of aggregate consumption, which is equal to productivity net of fixed operating costs, is very similar). As reported in panel B of Table 2, for our benchmark economy the total gains from trade are a 12.4 percent increase in aggregate productivity relative to autarky. This is, of course, an extreme comparison. In Table 3 we report the gains from trade for intermediate degrees of openness. In particular, holding all other parameters fixed, we change the trade cost τ so as to induce import shares of 0 (autarky), 10, 20, 30, and 38 percent (the Taiwan benchmark).

The model predicts a 3.4 percent increase in aggregate productivity moving from autarky to an import share of 10 percent. Moving to an import share of 20 percent adds another 2.7 percent so that the cumulative gain moving from autarky to an import share of 20 percent is $3.4 + 2.7 = 6.1$ percent. Continuing all the way to Taiwan's openness gives the 12.4 percent benchmark gains (relative to autarky) discussed above.

ACR show that, in a large class of models, the gains from trade are summarized by the formula $\frac{1}{\sigma}\log(\lambda/\lambda')$ where σ is the trade elasticity, as in (31) above, and where λ and λ' denote the aggregate share of spending on domestic goods before and after the change in trade costs. According to this formula, moving from autarky to an import share of 10 percent with a trade elasticity of 4.2 (which is what our model implies for that degree of openness) gives gains of $\frac{1}{4.2}\log(1/0.9) = 0.025$ or 2.5 percent. This is reasonably close to the 3.4 percent we find in our model. Similarly, according to this formula, moving from autarky to Taiwan's import share gives total gains of 11.7 percent, close to the 12.4 percent we find in our model. In short, even though our model with variable markups is not nested by the ACR setup, we find that their formula still provides a good approximation to the total gains from trade, especially for countries that are sufficiently away from autarky.

Intuitively, the ACR formula provides a good approximation even in our setting with variable markups because the trade elasticity itself endogenously captures important aspects of markup variation and consequently these aspects of markup variation are already reflected in the ACR gains. For example, when markups are not too variable, there is nearly one-for-one pass-through from changes in trade costs to changes in prices so the trade elasticity is high and the ACR gains are low. But when markups are highly variable, there is much less pass-through from changes in trade costs to changes in prices so the trade elasticity is low and the ACR gains are high.

B. *Procompetitive Gains from Trade*

We are now ready to ask the key question of this paper: to what extent does opening to international trade reduce product market distortions, i.e., the amount of misallocation induced by markups? This question has important implications for the design of policy reforms. In other words, we ask: do policymakers need to directly address product market distortions, or do they largely disappear if countries open up to trade? We argue below that opening to trade is a powerful substitute for much more complex, perhaps infeasible, product market policies aimed at reducing markup-induced misallocation.

We answer our question by studying the extent to which the losses from misallocation change as the economy opens up to trade. Notice that we can decompose TFP changes resulting from a trade policy into the change arising due to changes in the first-best level of TFP as well as due to the reduction in misallocation,

$$\Delta \log A = \Delta \log A_{\textit{efficient}} + \Delta(\log A - \log A_{\textit{efficient}}).$$

The first term on the right-hand side of this expression gives the change in the first-best level of TFP, while the second term gives reduction in misallocation, our measure of procompetitive effects. In a model with constant markups, aggregate productivity equals first-best productivity (the equilibrium allocation is efficient) and hence there are zero procompetitive gains. The procompetitive gains will be positive if increased trade reduces misallocation so that the increase in aggregate productivity is larger than the increase in first-best productivity. The procompetitive effects will be negative if increased trade increases misallocation.

Under autarky, the economy is 9 percent below the first-best level of productivity. With a 10 percent import share, the economy is 7.3 percent below the first-best. So, as reported in Table 3, the procompetitive gains from trade are 1.7 percent. Since misallocation is 9 percent in autarky, an import share of 10 percent reduces misallocation to $7.3/9 = 0.81$ of its autarky level, i.e., by almost 20 percent. Opening up all the way to Taiwan's import share gives somewhat larger procompetitive gains of 2 percent. Misallocation thus falls to $7.0/9.0$ or 0.78 of the level in autarky. Clearly, the extent of the reduction in misallocation, and hence the strength of the procompetitive effects, is largest near autarky and then diminishes in relative importance as the economy experiences increasing degrees of openness.

In short, opening up to trade reduces product market distortions in Taiwan by about one-fifth. As Table 2 reports, other measures of product market distortions, such as the dispersion in sector-level markups, fall in half under the Taiwan parameterization compared to autarky. International trade can thus significantly alleviate product market distortions.

C. *Further Discussion*

Domestic versus Import Markups.—As emphasized by Arkolakis et al. (2012), the overall size of the procompetitive effects depends on markup responses of producers both in their domestic market and in their export market. It can be the case that a reduction in trade barriers leads to lower domestic markups (as Home producers lose market share) combined with higher markups on imported goods (as Foreign producers gain market share), resulting in more misallocation—in which case the procompetitive "gains" from trade would be negative.[13] In short, looking only at the markups of domestic producers may be misleading. As reported in Table 3, we indeed see that markups on imported goods do increase as the economy opens to trade: the revenue-weighted harmonic average of markups on imported goods increases by 16.6 percent as the economy opens from autarky (where Foreign producers have infinitesimal market share) to an import share of 10 percent while the corresponding average for domestic (Home) markups falls by 1.6 percent. The latter fall receives much more weight in the economy-wide aggregate markup so that overall the aggregate markup falls 1.9 percent. Notice that the fall in the aggregate markup is larger than the fall in domestic markups alone. This is due to a compositional effect. In particular, although markups on imported goods are rising while domestic markups are falling, the *level* of domestic markups is higher than the level of markups on imported goods. As the economy opens, the aggregate

[13] See also Holmes, Hsu, and Lee (2014) who study misallocation in symmetric trade models with oligopolistic competition like ours and show that the effects of changes in trade costs on misallocation can be theoretically decomposed into two components, a cost-change component and a price-change component. While the cost-change component is always associated with a decrease in misallocation (i.e., is a source of positive procompetitive gains), the price-change component may lead to an increase in misallocation.

markup falls both because the high domestic markups of Home producers are falling and because a greater share of spending is on low-markup imports from Foreign producers.

Role of Cross-Country Correlation in Productivity.—To match an aggregate trade elasticity of $\sigma = 4$, our benchmark model requires a quite high degree of cross-country correlation in sectoral productivity draws, $\rho = 0.94$. This implies, that, following a reduction in trade barriers, there is a correspondingly high degree of *head-to-head competition* between producers within any given sector. In panel A of Table 4, we report the sensitivity of our results to the extent of correlation in sectoral productivity. For each level of ρ shown, we recalibrate our model to match our original targets except for the trade elasticity and related import share dispersion statistics. As we reduce ρ, the model trade elasticity falls monotonically, reaching values of less than 1. Corresponding to these low trade elasticities are extremely high total gains from trade. Mechanically, the trade elasticity falls because the index of import share dispersion $\mathrm{Var}[\lambda(s)]/\lambda(1-\lambda)$, i.e., the coefficient on θ in equation (33) above, *rises* as ρ falls. That is, an increasing proportion of sectors are either completely dominated by domestic producers (with import shares close to 0) or completely dominated by foreign producers (with import shares close to 1) so that the trade elasticity depends more on the across-sector θ and less on the within-sector elasticity γ.

When the correlation ρ is high, sectoral productivity draws are similar across countries so that most trade is *intraindustry*. In this case, a given change in trade costs gives rise to relatively large changes in trade flows. Panel A of Table 4 shows that the Grubel and Lloyd (1971) index of intraindustry trade is monotonically decreasing in ρ, falling from 0.45 for our benchmark model (meaning, 45 percent of trade is intraindustry) to less than 0.1 for $\rho < 0.5$. We also note that in our benchmark model there is a strong positive relationship between a sector's share of domestic sales and its share of imports. In particular, the slope coefficient in a regression of sector imports as a share of total imports on sector domestic sales as a share of total domestic sales is about 0.55—i.e., sectors with large, productive firms are also sectors with large import shares, which suggests firms in these sectors face a great deal of head-to-head competition. When we reduce ρ we find this regression coefficient falls, eventually becoming slightly negative, so that large sectors no longer have large import shares, suggesting domestic producers no longer face as much competition when ρ is low.

Importantly, when the correlation ρ is sufficiently low a reduction in trade costs actually increases misallocation so that, as in Arkolakis et al. (2012), the procompetitive "gains" from trade are negative. To see this, begin with an economy with high correlation, $\rho = 0.9$ (similar to our benchmark). As shown in panel B of Table 4, there is a substantial fall in markup dispersion across domestic producers. Ultimately this is a consequence of hitherto dominant domestic producers losing substantial market share to foreign competition. With $\rho = 0.9$, opening to trade reduces misallocation from 9 percent to 7.1 percent so that there are positive procompetitive gains of $9 - 7.1 = 1.9$ percent. By contrast, with less correlation in draws, say $\rho = 0.1$, opening to trade increases misallocation from 9 percent to 10.9 percent so that there are *negative* procompetitive "gains" of $9 - 10.9 = -1.9$ percent. In this case, misallocation is worse with trade than it is under autarky. This subtracts

TABLE 4—IMPORTANCE OF HEAD-TO-HEAD COMPETITION

Panel A. Sensitivity to cross-country correlation, ρ

					Gains from trade, percent		
ρ	σ	Imp. share dispersion	Intra-industry	Imp. share on sales	Pro-C.	Total	ACR
1.00	4.41	0.17	0.56	0.66	2.1	12.0	10.7
0.90	3.61	0.34	0.38	0.48	1.9	13.2	13.0
0.80	2.68	0.51	0.25	0.36	1.5	16.7	17.5
0.70	2.14	0.62	0.18	0.29	1.2	21.3	22.0
0.60	1.76	0.68	0.13	0.26	0.9	26.0	26.7
0.50	1.50	0.74	0.10	0.24	0.6	31.2	31.3
0.40	1.30	0.78	0.07	0.22	0.2	36.7	36.1
0.30	1.14	0.81	0.05	0.19	−0.2	42.7	41.4
0.20	0.99	0.84	0.04	0.15	−0.9	49.9	47.4
0.10	0.85	0.87	0.03	0.06	−1.9	58.6	55.1
0.00	0.66	0.91	0.02	−0.09	−3.5	62.8	71.8
0.94	4.00	0.26	0.45	0.55	2.0	12.4	11.7
data	4.00	0.38	0.37	0.81			

Panel B. Markup dispersion and cross-country correlation

	ρ = 0.9		ρ = 0.1	
	Autarky	Taiwan	Autarky	Taiwan
Misallocation				
TFP loss, percent	9.0	7.1	9.0	10.9
All markups				
Aggregate markup	1.35	1.32	1.35	1.40
SD log	0.090	0.083	0.090	0.092
log p99/p50	0.390	0.417	0.390	0.417
Domestic markups				
Aggregate markup	1.35	1.32	1.35	1.39
SD log	0.090	0.074	0.090	0.097
log p99/p50	0.390	0.376	0.390	0.439
Import markups				
Aggregate markup	1.11	1.32	1.11	1.40
SD log	0	0.104	0	0.085
log p99/p50	0	0.486	0	0.385

Notes: Panel A shows the sensitivity of our results to the cross-country correlation ρ in sectoral productivity. For each ρ we recalibrate the model targeting our usual moments except for the trade elasticity σ, the index of import share dispersion, the Grubel-Lloyd index of intraindustry trade, and the slope coefficient in a regression of sectoral import shares on sectoral sales share. The last three columns show the procompetitive gains from trade, total gains from trade, and gains predicted by the ACR formula. The second-last row shows the equivalent results for our benchmark calibration. Panel B shows the implications for markup dispersion under high $\rho = 0.9$ and low $\rho = 0.1$ levels of correlation.

from the total gains from trade (which are nonetheless very large here, because of the counterfactually low trade elasticity with $\rho = 0.1$).

In panel A of Table 4, we also report the data counterparts of the index of import share dispersion, the Grubel and Lloyd index, and the coefficient of size on import shares. To match these, our model requires ρ in the range 0.8 to 1.0 (depending on how much weight is given to each measure) with the trade elasticity then being in the range 2.7 to 4.4. In short, to match the import share dispersion, intraindustry trade, and the trade elasticity, the model requires a high degree of cross-country correlation in productivity draws.

Finally, in panel A of Table 4, we also report the total gains from trade implied by the ACR formula for the values of the trade elasticity σ shown. Notice that while the

ACR formula provides a good approximation to the total gains (especially when the trade elasticity is not too low), the decomposition of those gains into procompetitive and other channels depends quite sensitively on the parameterization.

Alternative Model: Cross-Country Correlation in Idiosyncratic Draws.—As a final way to see the importance of head-to-head competition, we provide results for an alternative version of our model where there is correlation in *both* sectoral productivities $z(s)$ and in producer-specific idiosyncratic draws $x_i(s)$. Specifically we assume $H_Z(z, z^*) = \mathcal{C}_Z(F_Z(z), F_Z(z^*))$ and $H_X(x, x^*) = \mathcal{C}_X(F_X(x), F_X(x^*))$ both linked via a Gumbel copula as in (37) but with distinct correlation coefficients, ρ_z and ρ_x. The benchmark model is then the special case $\rho_z = 0.94$ and $\rho_x = 0$. We recalibrate this model targeting the same moments as our benchmark model plus one new moment that helps identify ρ_x. In particular, we choose ρ_x so that our model reproduces the cross-sectional relationship between sectoral *import penetration* and sectoral *concentration amongst domestic producers* that we observe in the Taiwanese data. In the data, the slope coefficient in a regression of sector import penetration on sector domestic HH indexes is 0.21—i.e., sectors with high import penetration are also sectors with relatively high concentration amongst domestic producers.[14] To match this, we require a slight degree of correlation in idiosyncratic draws, $\rho_x = 0.05$. The required correlation in sectoral productivity is correspondingly slightly lower, $\rho_z = 0.93$.

As reported in panel B of Table 3, this alternative model implies very similar total gains from trade, 12 percent versus the benchmark 12.4 percent, but because dominant producers face more head-to-head competition there are now larger procompetitive effects. Opening from autarky to Taiwan's import share now reduces misallocation by almost one-third and the procompetitive gains are 2.4 percent, up from the benchmark 2 percent. Here, trade plays a larger role in reducing markup distortions because countries import more of exactly those goods for which the domestic market is in fact more distorted.

Capital Accumulation and Elastic Labor Supply.—In the benchmark model the only gains are from changes in aggregate productivity. The aggregate markup falls 2.9 percent between autarky and the Taiwan benchmark but this change in the aggregate markup has no welfare implications. In contrast, with capital accumulation and/or elastic labor supply the aggregate markup acts like a distortionary wedge affecting investment and labor supply decisions, and, because of this, a reduction in the aggregate markup also increases welfare. In particular, suppose the representative consumer has intertemporal preferences $\sum_{t=0}^{\infty} \beta^t U(C_t, L_t)$ over aggregate consumption C_t and labor L_t and that capital is accumulated according to $K_{t+1} = (1 - \delta)K_t + I_t$. Suppose also that individual producers have production function $y = ak^{\alpha} l^{1-\alpha}$. We solve this version of the model assuming a utility function $U(C, L) = \log C - L^{1+\eta}/1 + \eta$, a discount factor $\beta = 0.96$, a depreciation rate $\delta = 0.1$, an output elasticity of capital $\alpha = 1/3$, and several alternative values for the elasticity of labor supply η. We start the economy in autarky and then

[14] For our benchmark model, the slope coefficient of sector import penetration on sector domestic HH indexes is 0.14, somewhat low relative to the 0.21 in the data. See the online Appendix for more details.

TABLE 5—GAINS FROM TRADE WITH ELASTIC FACTORS

		Variable markups		
		Frisch elasticity of labor supply $(1/\eta)$		
	Constant markups	0	1	∞
Change TFP, percent	10.4	12.4	12.4	12.4
Change markup, percent	0	−2.9	−2.9	−2.9
Change C, percent	15.7	19.5	21.3	23.0
Change K, percent	15.7	23.0	24.8	26.6
Change Y, percent	15.7	20.1	21.9	23.7
Change L, percent	0	0	1.8	3.6
Change welfare, percent (including transition)	14.5	18.0	18.1	18.4
Procompetitive gains	0	3.5	3.6	3.9

Notes: Representative consumer has preferences $\sum_{t=0}^{\infty} \beta^t U(C_t, L_t)$ over aggregate consumption C_t and labor L_t with $U(C, L) = \log C - L^{1+\eta}/1 + \eta$. Capital is accumulated according to $K_{t+1} = (1-\delta)K_t + I_t$. Individual producers have production function $y = ak^{\alpha} l^{1-\alpha}$. We set discount factor $\beta = 0.96$, depreciation rate $\delta = 0.1$, output elasticity of capital $\alpha = 1/3$ and elasticities of labor supply η as shown.

compute the transition to a new steady-state corresponding to the Taiwan benchmark. We measure the welfare gains as the consumption compensating variation taking into account the dynamics of consumption and employment during the transition to the new steady-state. As reported in Table 5, with capital accumulation and a Frisch elasticity of 1, the welfare gains are 18.1 percent of which 3.6 percent is due to procompetitive effects. These are about 1.5 times larger than in our benchmark setup with inelastic factors in which welfare increases by 12.4 percent of which 2 percent is due to procompetitive effects.

V. Robustness Experiments

We now consider variations of our benchmark model, each designed to examine the sensitivity of our results to parameter choices or other assumptions. For each robustness experiment we recalibrate the trade cost τ, export fixed cost f_x, and correlation parameter ρ so that the Home country continues to have an aggregate import share of 0.38, fraction of exporters 0.25, and trade elasticity 4, as in our benchmark model. A summary of these robustness experiments is given in Table 6. Further details and a full set of results for these experiments are reported in the online Appendix.

Heterogeneous Labor Market Distortions.—Our benchmark model focuses on the importance of product market distortions but ignores the role of labor market distortions. We now show that this is not essential for our main results. We assume that there is a distribution of producer-level labor market distortions that act like labor input taxes, putting a *wedge* between labor's marginal product and its factor cost. Specifically, a producer with productivity a also faces an input tax $t(a)$ on its wage bill so that it pays $(1 + t(a))W$ for each unit of labor hired. We assume $t(a) = \frac{a\tau_l}{1 + a\tau_l}$ and choose the parameter τ_l governing the sensitivity of the labor distortion to producer productivity so that our model matches the spread between the average producer labor share and the aggregate labor share that we observe in

Table 6—Robustness Experiments

	Benchmark	Labor wedge	Tariffs	Bertrand	Low γ	High γ	No fixed	Gaussian
Trade elasticity	4.00	4.00	4.00	4.00	3.59	4.00	4.00	4.00
Import share	0.38	0.38	0.38	0.38	0.38	0.38	0.38	0.38
Fraction exporters	0.25	0.25	0.25	0.25	0.25	0.25	1	0.25
TFP loss autarky	9.0	8.8	9.0	4.6	4.9	11.3	8.9	9.8
TFP loss Taiwan	7.0	6.8	6.9	2.1	2.3	9.9	6.9	7.2
Gains from trade	12.4	12.2	14.6	13.8	16.6	11.8	11.8	11.6
Procompetitive gains	2.0	2.0	2.0	2.5	2.7	1.4	2.0	2.6
Key parameters								
ρ	0.94	0.93	0.94	0.91	1.00	0.92	0.95	0.99
τ	1.129	1.129	1.055	1.132	1.080	1.138	1.136	1.129
f_x	0.203	0.065	0.055	0.110	1.350	0.040	0	0.070
Additional moments								
Avg./agg. labor share	1.16	1.35	1.17	1.09	1.07	1.23	1.19	1.20
Mean tariff			0.062					
SD tariff			0.039					

the data. In the data, the average producer labor share is 1.35 times the aggregate labor share. Since the latter is a weighted version of the former, this tells us that large producers tend to have low measured labor shares. To match this, our model requires $\tau_l = 0.001$, so indeed producers with high productivities are also producers with relatively high labor distortions but the correlation is quite weak.[15]

These labor market distortions significantly reduce aggregate productivity relative to the benchmark economy—the level of productivity turns out to be only about 80 percent of the benchmark. In this sense, total misallocation is much larger in this economy. But this is because there are now *two* sources of misallocation—labor market distortions and markup distortions. The amount of misallocation due to markup distortions alone is roughly the same as in the benchmark economy. To see this, notice that the level of productivity associated with a planner who faces the same labor distortions but can otherwise reallocate across producers is 6.8 percent higher than the equilibrium level of productivity, very close to the corresponding 7 percent gap in the benchmark economy.

Given that there are similar amounts of misallocation due to markups, it is not then surprising that the gains from trade turn out to be similar as well. The aggregate gains from trade are 12.2 percent versus the benchmark 12.4 percent while the procompetitive gains are about 2 percent in both cases. In short, allowing for labor market distortions does not change our results.

Heterogeneous Tariffs.—In our benchmark model, the only barriers to trade are the iceberg trade costs τ and the fixed cost of exporting f_x and these are the same for every producer in every sector. We consider a version of our model where in addition to these trade costs there is a *sector-specific* distortionary tariff that is levied on the value of imported goods. For simplicity we assume that tariff revenues are

[15] For our benchmark model the average labor share is also greater than the aggregate labor share, but the spread is 1.16, somewhat lower than the 1.35 in the data.

rebated lump-sum to the representative consumer. We assume the tariff rates are drawn from a beta distribution on [0, 1] with parameters estimated by maximum likelihood using the Taiwanese micro data. These estimates imply a mean tariff rate of 0.062 with cross-sectional standard deviation of 0.039. With a mean tariff of 0.062, the trade cost required to match the aggregate import share is correspondingly lower, 1.055 down from the benchmark 1.129.

Perhaps surprisingly, we find the total gains from trade are somewhat larger than in the benchmark, 14.6 percent as opposed to 12.4 percent, with the procompetitive gains 2 percent, the same as in the benchmark. One might expect that, for a given distribution of tariffs, a symmetric reduction in trade costs would make the cross-sectoral misallocation due to tariffs worse and thereby reduce the gains from trade (relative to an economy without tariffs). Instead, we find that there is a substantial "second best" effect—i.e., in the presence of two distortions, increasing one distortion does not necessarily reduce welfare. In particular, the additional cross-sectoral misallocation due to tariffs is offset by strong reductions in within-sector market share dispersion.

Bertrand Competition.—In our benchmark model, firms engage in Cournot competition. If we assume instead that firms engage in Bertrand competition, then the model changes in only one respect. The demand elasticity facing producer i in sector s is no longer a harmonic weighted average of θ and γ, as in equation (13), but is now an arithmetic weighted average, $\varepsilon_i(s) = \omega_i(s)\theta + \big(1 - \omega_i(s)\big)\gamma$. With this specification the results are similar to the benchmark. The total gains from trade are 13.8 percent, up from the benchmark 12.4 percent, and the procompetitive gains are 2.5 percent, likewise up slightly from the benchmark 2 percent. As shown in the last two columns of Table 2, the Bertrand model implies somewhat lower markup dispersion than the Cournot model. But it also implies a larger *change* in markup dispersion when opening to trade and hence a larger reduction in misallocation. Opening from autarky to Taiwan's import share implies that misallocation falls by one-half, up from the benchmark one-fifth fall. Perhaps not surprisingly, the competitive pressure on dominant firms following a trade liberalization is greater with Bertrand competition than with Cournot. Consequently, the Bertrand model implies, if anything, larger procompetitive effects than the benchmark.

Role of γ.—For our benchmark calibration procedure we obtain $\gamma = 10.5$, quite close to the value $\gamma = 10$ used by Atkeson and Burstein (2008). To see what features of the data determine this value, we have recalibrated our model with a range of alternate values for γ.

For example, with a much lower value of $\gamma = 5$ we find that the model cannot produce a trade elasticity of 4—even setting $\rho = 1$ (perfect correlation) gives a trade elasticity of only 3.59. In addition, as reported in the online Appendix, with $\gamma = 5$ the model also implies too much intraindustry trade, too little import share dispersion, and too strong an association between sector import shares and size. At the other extreme, with a much higher value of $\gamma = 20$, the model can better match the trade elasticity and facts on intraindustry trade and import share dispersion, but now produces too weak an association between sector import shares and size as well as too strong an association between sector concentration and import penetration.

In short, low values like $\gamma = 5$ and high values like $\gamma = 20$ are both at odds with key features of the data. In trying to match these moments, our calibration procedure selects the value $\gamma = 10.5$. Importantly, our model's implications for the gains from trade do not change dramatically even for these extreme values of γ. For example, with $\gamma = 5$ the model implies that the total gains are 16.6 percent of which 2.7 percent are procompetitive gains. With $\gamma = 20$ the model implies that total gains are 11.8 percent of which 1.4 percent are procompetitive gains.

No Fixed Costs.—To assess the role of the fixed costs f_d and f_x we compute results for a version of our model with $f_d = f_x = 0$. In this specification, all firms operate in both their domestic and export markets. Hence the equilibrium number of producers in a sector is simply pinned down by the geometric distribution for $n(s)$. This version of the model yields almost identical results to the benchmark. Shutting down these extensive margins makes little difference because the typical producer near the margin of operating or not is small and has negligible impact on the aggregate outcomes.

Gaussian Copula.—Our benchmark model uses the Gumbel copula (37) to model cross-country correlation in sectoral productivity draws. To examine the sensitivity of our results to this functional form, we resolve our model using a Gaussian copula, namely

$$\mathcal{C}(u, u^*) = \Phi_{2,r}(\Phi^{-1}(u), \Phi^{-1}(u^*)), \tag{40}$$

where $\Phi(x)$ denotes the CDF of the standard normal distribution and $\Phi_{2,r}(x, x^*)$ denotes the standard bivariate normal distribution with linear correlation coefficient $r \in (-1, 1)$. To compare results to the Gumbel case, we map the linear correlation coefficient into our preferred Kendall correlation coefficient, which for the Gaussian copula is $\rho = 2\arcsin(r)/\pi$. To match a trade elasticity of 4 requires $\rho = 0.99$, up from the benchmark 0.94 value. This version of the model also yields very similar results to the benchmark. Conditional on choosing the amount of correlation to match the trade elasticity, the total gains from trade are 11.6 percent with procompetitive gains of 2.6 percent, both quite close to their benchmark values. In short, our results are not sensitive to the assumed functional form of the copula.

VI. Extensions

A. *Asymmetric Countries*

Our benchmark model makes the stark simplifying assumption of trade between two symmetric countries. We now relax this and consider the gains from trade between countries that differ in size and/or productivity. Specifically, we normalize the Home country labor force to $L = 1$ and vary the Foreign labor force L^*. Home producers continue to have production function $y_i(s) = a_i(s)l_i(s)$, as in (3) above, and Foreign producers now have the production function $y_i^*(s) = \overline{A}^* a_i^*(s)l_i^*(s)$ with productivity scale parameter $\overline{A}^*$. We again recalibrate key parameters of the model so that for the Home country we reproduce the degree of openness of the

TABLE 7—GAINS FROM TRADE WITH ASYMMETRIC COUNTRIES

	Benchmark		*Panel A. Larger trading partner*			
			$L^* = 2L$		$L^* = 10L$	
	Home	Foreign	Home	Foreign	Home	Foreign
ρ	0.94	0.94	0.94	0.94	0.96	0.96
Trade elasticity	4.00	4.00	4.00	4.19	4.00	4.30
Import share	0.38	0.38	0.38	0.21	0.38	0.05
Fraction exporters	0.25	0.25	0.25	0.11	0.25	0.05
TFP loss autarky	9.0	9.0	9.0	9.0	9.0	9.0
TFP loss	7.0	7.0	7.3	6.8	7.1	7.5
Gains from trade	12.4	12.4	12.0	6.5	12.0	2.2
Procompetitive gains	2.0	2.0	1.7	2.1	1.9	1.4
			Panel B. More productive trading partner			
			$\overline{A}^* = 2\overline{A}$		$\overline{A}^* = 10\overline{A}$	
			Home	Foreign	Home	Foreign
ρ			0.86	0.86	0.60	0.60
Trade elasticity			4.00	3.12	4.00	1.30
Import share			0.38	0.21	0.38	0.07
Fraction exporters			0.25	0.10	0.25	0.03
TFP loss autarky			9.0	9.0	9.0	9.0
TFP loss			7.4	7.0	7.5	8.6
Gains from trade			15.3	8.1	31.9	5.6
Procompetitive gains			1.5	1.8	1.5	0.4

Taiwan benchmark—in particular, we choose the proportional trade cost τ, export fixed cost f_x, and correlation parameter ρ so that the Home country continues to have an aggregate import share of 0.38, fraction of exporters 0.25 and trade elasticity $\sigma = 4$.

Larger Trading Partner.—The top panel of Table 7 shows the gains from trade when the Foreign country has labor force $L^* = 2$ and $L^* = 10$ times as large as the Home country. For the Home country, the total gains from trade are slightly smaller than under symmetry. And when the Foreign country is larger, its total gains from trade are smaller than the Home country gains. For example, when the Foreign country is ten times as large as the Home country, the Home gains are 12 percent (down from 12.4 percent in the symmetric benchmark) whereas the Foreign gains are down to 2.2 percent. The Home country has much more to gain from integration with a large trading partner than the Foreign country has to gain from integration with a small trading partner. The procompetitive gains are also slightly lower for both countries. When $L^* = 10$, the Home procompetitive gains are 1.9 percent (down from 2 percent in the symmetric benchmark) whereas the Foreign procompetitive gains are down to 1.4 percent. Interestingly, the procompetitive gains account for a high share of the Foreign country's total gains, 1.4 percent out of 2.2 percent. In this calibration, the Foreign country is considerably less open than the Home country, with an aggregate import share of 0.05 (as opposed to 0.38) and a fraction of exporters of 0.05 (as opposed to 0.25). Despite the lower openness, we see that Foreign consumers still gain considerably from exposing their producers to greater competition (Home consumers gain even more), and that failing to account for procompetitive effects can seriously understate the gains from integration, even for a large country.

More Productive Trading Partner.—The bottom panel of Table 7 shows the gains from trade when the Foreign country has productivity scale $\overline{A}^* = 2$ and $\overline{A}^* = 10$ times that of the Home country but has the same size, $L^* = 1$. Not surprisingly, for the Home country the total gains from trade are considerably larger than under symmetry. For example, when $\overline{A}^* = 10$, the Home gains are 31.9 percent (up from 12.4 percent in the symmetric benchmark). But these very large gains are almost entirely due to increases in the first-best level of productivity. The procompetitive gains are 1.5 percent, and hence relative to the symmetric benchmark are both smaller in absolute terms and smaller as a share of the total gains. The more productive Foreign country has smaller total gains (and so benefits less from trade than the less productive Home country) and smaller procompetitive gains.

The correlation in cross-country productivity required to reproduce a Home trade elasticity of $\sigma = 4$ is $\rho = 0.60$, considerably lower than the benchmark $\rho = 0.94$. With large productivity differences between countries, import shares are more responsive to changes in trade costs than under symmetry. But because there is less correlation, there is also less head-to-head competition and so the procompetitive gains are smaller.

B. *Free Entry*

In our benchmark model there is an exogenous number $n(s)$ of firms in each sector, a subset of which choose to pay the fixed cost f_d and operate. Some of the firms that do operate make substantial economic profits and thus there is an incentive for other firms to try to enter. We now relax the no-entry assumption and assume instead that there is *free entry* subject to a sunk cost. In equilibrium, the expected profits simply compensate for this initial sunk cost.

To keep the analysis tractable, we assume that entry is not directed at a particular sector. After paying its sunk cost, a firm learns the productivity with which it operates, as in Melitz (2003), as well as the sector to which it is randomly assigned.[16] We also assume that there are no fixed costs of operating or exporting in any given period. Instead, we assume that a firm's productivity is drawn from a discrete distribution which includes a mass point at zero, thus allowing the model to generate dispersion in the number of firms that operate.

Computational Issues.—Given the structure of our model, the *expected* profits of a potential entrant (which, due to free entry, equals the sunk cost) are not equal to the *average* profits across those firms that operate. One reason for this difference is that a potential entrant recognizes the effect its entry will have on its own profits and those of the incumbents. An additional reason is that the measure of producers of different productivities in a given sector is correlated with the profits a particular firm makes in that sector. Computing the expected profits of a potential entrant is thus a challenging task: we need to integrate the distribution (across sectors) of the measures of firms (over their productivities)—a finite, but high-dimensional object. In addition, a potential entrant must resolve for the distribution of markups

[16] An unappealing implication of allowing *directed entry* is that the resulting model would predict low dispersion in sectoral markups, in stark contrast to the very high dispersion in sectoral markups in the data.

that would arise if it enters. Given that the number of firms that enter each sector is small, the law of large number fails, and the algorithm to compute an equilibrium is involved. For this reason, we make a number of additional simplifying assumptions relative to our benchmark model without entry. In particular, we use a coarse productivity distribution and set the operating and exporting fixed costs to $f_d = f_x = 0$.

Setup.—The productivity of a firm in sector $s \in [0, 1]$ is now given by a world component, common to both countries, $z(s)$, and a firm-specific component. In addition, we assume a *gap* $u(s)$ between the productivity with which firms produce for their domestic market and that with which they produce for their export market. Specifically, let $u(s)$ denote the productivity gap for Home producers in sector s and let $u^*(s)$ denote the productivity gap for Foreign producers in sector s. There is an unlimited number of potential entrants. To enter, a firm pays a sunk cost f_e that allows it to draw (i) a sector s in which to operate, and (ii) idiosyncratic productivity $x_i(s) \in \{0, 1, \bar{x}\}$. To summarize, a Home firm in sector s with idiosyncratic productivity $x_i(s)$ produces for its domestic market with overall productivity $a_i^H(s) = z(s)u(s)x_i(s)$ and produces for its export market with overall productivity $a_i^{*H}(s) = z(s)x_i(s)/\tau$ where τ is the gross trade cost. Similarly, a Foreign firm in sector s with idiosyncratic productivity $x_i^*(s)$ produces for its domestic market with overall labor productivity $a_i^{*F}(s) = z(s)u^*(s)x_i^*(s)$ and produces for its export market with overall productivity $a_i^F(s) = z(s)x_i^*(s)/\tau$.

Cross-Country Correlation and Head-to-Head Competition.—In this version of the model, the amount of head-to-head competition can now be varied flexibly by changing the amount of dispersion in $u(s)$ across sectors. Greater dispersion in $u(s)$ reduces the amount of head-to-head competition between Home and Foreign producers and thereby lowers the aggregate trade elasticity.

Parameterization.—The Taiwanese data feature a high degree of across-sector dispersion in markups, in the number of producers, and in market concentration. We match this across-sector dispersion by assuming that the probability that a firm draws idiosyncratic productivity $x_i(s) \in \{0, 1, \bar{x}\}$ varies with s (but is the same across countries for a given sector). In particular, we assume a nonparametric distribution $\text{Prob}[x_i(s) \mid s]$ across sectors and calibrate this distribution to match the same set of moments we targeted for our benchmark model (we have found that allowing for nine types of sectors produces a good fit; in the online Appendix we also discuss results for a simpler model with a single sector type).

We assume that the gaps $u(s)$ are drawn from a log-normal distribution with variance σ_u^2 and that the worldwide sectoral productivities $z(s)$ are drawn from a Pareto distribution with shape parameter ξ_z.

Taiwan Calibration Revisited.—We fix $\gamma = 10.5$ and $\theta = 1.24$, as in our benchmark model. We calibrate the new parameters f_e, $\bar{x}$, σ_u, ξ_z, the distribution $\text{Prob}[x_i(s) \mid s]$ across sectors, and the trade cost τ targeting the same moments as in our benchmark model. The full set of results for this calibration are reported in the online Appendix.

Table 8—Entry and Collusion

	Panel A. No collusion			Panel B. 25 percent collusion		
	No entry	Free entry	Autarky	No entry	Free entry	Autarky
Number of firms trying to enter	187	168	187	162	160	162
TFP loss	2.0	2.2	3.4	4.6	4.6	9.0
Total fixed costs	0.45	0.41	0.45	0.31	0.30	0.31
Aggregate profits	0.45	0.46	0.51	0.38	0.38	0.41
Aggregate markup	1.25	1.26	1.32	1.27	1.27	1.35
Expected profits, entrants	0.22	0.24	0.24	0.19	0.19	0.19
Gains from trade	8.2	6.9		11.7	11.6	
Procompetitive gains	1.5	1.2		4.4	4.3	
Misallocation relative to autarky	0.57	0.64		0.51	0.52	
Change markup, percent	−5.8	−5.0		−6.3	−6.3	

Gains from Trade with Free Entry.—Panel A of Table 8 shows the gains from trade in this economy. With free entry, 168 firms pay the sunk cost and enter any individual sector. The economy is about 2.2 percent away from the first-best level of aggregate productivity. Thus, although we target the same concentration moments and have the same elasticities θ and γ as in the benchmark model, with free entry there is less misallocation.

Aggregate productivity is 6.9 percent above its autarky level and opening to trade reduces misallocation by just over one-third, from 3.4 percent to 2.2 percent. This reduction in misallocation implies procompetitive gains of 1.2 percent, somewhat lower than the benchmark procompetitive gains of 2 percent. Note that there are 187 firms attempting to enter under autarky, more than in the open economy. For a given number of firms, expected profits are higher under autarky and so more firms enter until the free-entry condition is satisfied. If we hold the number of firms fixed at the autarky level of 187 but otherwise open the economy to trade, aggregate productivity rises by 8.2 percent, larger than the 6.9 percent with free entry, and the procompetitive gains are correspondingly larger at 1.5 percent as opposed to 1.2 percent.

To summarize, even with free entry there is a quantitatively significant reduction in misallocation. Importantly, the somewhat weaker procompetitive effects reflect the alternative calibration of the model which implies less initial misallocation, not the free entry itself. In particular, the model predicts much less dispersion in sectoral markups—e.g., the ratio of the ninetieth percentile to the median is 1.14 (compared to 1.27 in the data and 1.24 in the benchmark), and the ratio of the ninety-fifth percentile to the median is 1.17 (1.50 in the data and 1.42 in the benchmark).[17] We address this discrepancy between the model and the data next.

Collusion.—Given this failure to match the dispersion of sectoral markups in the data, we now consider a slight variation on the free-entry model designed to bridge the gap between the model and the data along this dimension. We suppose that with

[17] The online Appendix reports results for these experiments in more detail.

probability ψ all the high-productivity firms (those with $x_i(s) = \bar{x} > 1$) within a given sector are able to *collude*.[18] These colluding firms choose a single price to maximize their *group profits*. Since their collective market share is larger than their individual market shares, the price set by colluding firms is higher than the price they would charge in isolation and hence their collective markup is also correspondingly larger. Since this version of the model features more dispersion in markups, it also features more misallocation.

Panel B of Table 8 shows results for this model with $\psi = 0.25$. Even with free entry, this version of the model features productivity losses of 4.6 percent relative to the first-best. The reason these productivity losses are greater is that now the dispersion in sectoral markups is greater. For example, the ratio of the ninetieth percentile to the median is 1.22 (compared to 1.14 absent collusion) and the ratio of the ninety-fifth percentile to the median is 1.31 (compared to 1.17 absent collusion). Thus this version of the model produces sectoral dispersion in markups more in line with our benchmark model and hence closer to the data.

Consequently, the model now predicts larger total gains from trade of 11.6 percent, of which 4.3 percent are procompetitive gains—i.e., the model with free entry and collusion implies larger procompetitive gains than our benchmark model. With widespread collusion amongst domestic producers, opening to foreign competition provides an import source of market discipline. Notice also that the number of producers change very little (from 162 in autarky to 160 in the open economy) despite the reduction in firm markups (the aggregate markup falls from 1.35 to 1.27). The reason the number of firms does not change much is an externality akin to that in Blanchard and Kiyotaki (1987). Although an individual firm loses profits if its own markup falls, it benefits when other firms reduce markups due to the increase in aggregate output and the reduction in the aggregate price level. Overall, these two effects on expected profits roughly cancel each other out so that there is little effect on the gains from trade.

In short, with free entry and collusion the model implies strong procompetitive effects. In the online Appendix we report results for a wide range of collusion probabilities ψ and show that the same basic pattern holds. For example, if the collusion probability is $\psi = 0.15$ instead of $\psi = 0.25$ then the total gains from trade are 12.5 percent of which 4.2 percent are procompetitive gains.

The results from the model with collusion reinforce our main message: the procompetitive gains from trade are larger when product market distortions are large to begin with.

VII. Conclusions

We study the procompetitive gains from international trade in a quantitative model with endogenously variable markups. We find that trade can significantly reduce markup distortions if two conditions are satisfied: (i) there must be large inefficiencies associated with markups, i.e., extensive misallocation, and (ii) trade must in fact expose producers to greater competitive pressure. The second condition

[18] Alternatively, this can be thought of as the result of mergers or acquisitions.

is satisfied if there is a high degree of cross-country correlation in the productivity with which producers within a given sector operate.

We calibrate our model using Taiwanese producer-level data and find that these two conditions are satisfied. The Taiwanese data is characterized by a large amount of dispersion and concentration in producer market shares and a strong cross-sectional relationship between producer market shares and markups, which implies extensive misallocation. Moreover to match standard estimates of the trade elasticity, and at the same time match key facts on import share dispersion, intraindustry trade, and the cross-sectional relationship between import penetration and domestic concentration, the model requires a high degree of cross-country correlation in productivity. Consequently, the model implies that opening to trade does in fact expose producers to considerably greater competitive pressure.

We find that opening to trade reduces misallocation by about one-fifth in our benchmark model with Cournot competition and by about one-half in our model with Bertrand competition. Likewise, in our alternate model with free entry, opening to trade reduces misallocation by between one-third and one-half, depending on the specification. In this sense, we find that, indeed, trade can significantly reduce product market distortions.

We conclude by noting that, from a policy viewpoint, our model suggests that obtaining large welfare gains from an improved allocation of resources may not require a detailed, perhaps impractical, scheme of producer-specific subsidies and taxes that reduce the distortions associated with variable markups. Instead, simply opening an economy to trade may provide an excellent practical alternative that substantially improves productivity and welfare. Conversely, our model also predicts that countries which open up to trade after having already implemented policies aimed at reducing markup distortions may benefit less from trade than countries with large product market distortions. The former countries would mostly receive the standard gains from trade, while the latter would also benefit from the reduction in markup distortions.

REFERENCES

Arkolakis, Costas, Arnaud Costinot, David Donaldson, and Andrés Rodríguez-Clare. 2015. "The Elusive Pro-Competitive Effects of Trade." National Bureau of Economic Research Working Paper 21370.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *American Economic Review* 102 (1): 94–130.

Atkeson, Andrew, and Ariel Burstein. 2008. "Pricing-to-Market, Trade Costs, and International Relative Prices." *American Economic Review* 98 (5): 1998–2031.

Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum. 2003. "Plants and Productivity in International Trade." *American Economic Review* 93 (4): 1268–90.

Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott. 2010. "Multiple-Product Firms and Product Switching." *American Economic Review* 100 (1): 70–97.

Blanchard, Olivier Jean, and Nobuhiro Kiyotaki. 1987. "Monopolistic Competition and the Effects of Aggregate Demand." *American Economic Review* 77 (4): 647–66.

Chaney, Thomas. 2008. "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *American Economic Review* 98 (4): 1707–21.

De Blas, Beatriz, and Katheryn Russ. 2015. "Understanding Markups in the Open Economy." *American Economic Journal: Macroeconomics* 7 (2): 157–80.

De Loecker, Jan, Pinelopi Goldberg, Amit Khandelwal, and Nina Pavcnik. 2014. "Prices, Markups and Trade Reform." http://www.princeton.edu/~jdeloeck/india_dLGKP-1.pdf (accessed August 17, 2015).

De Loecker, Jan, and Frederic Warzynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review* 102 (6): 2437–71.

Devereux, Michael B., and Khang Min Lee. 2001. "Dynamic Gains from International Trade with Imperfect Competition and Market Power." *Review of Development Economics* 5 (2): 239–55.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70 (5): 1741–79.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2015. "Competition, Markups, and the Gains from International Trade: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.20120549.

Epifani, Paolo, and Gino Gancia. 2011. "Trade, Markup Heterogeneity and Misallocations." *Journal of International Economics* 83 (1): 1–13.

Feenstra, Robert C. 2003. "A Homothetic Utility Function for Monopolistic Competition Models, Without Constant Price Elasticity." *Economics Letters* 78 (1): 79–86.

Feenstra, Robert C., and David E. Weinstein. 2010. "Globalization, Markups, and the U.S. Price Level." National Bureau of Economic Research Working Paper 15749.

Grubel, Herbert G., and P. J. Lloyd. 1971. "The Empirical Measurement of Intra-Industry Trade." *Economic Record* 47 (120): 494–517.

Harrison, Ann E. 1994. "Productivity, Imperfect Competition and Trade Reform: Theory and Evidence." *Journal of International Economics* 36 (1-2): 53–73.

Holmes, Thomas J., Wen-Tai Hsu, and Sanghoon Lee. 2014. "Allocative Efficiency, Mark-Ups, and the Welfare Gains from Trade." *Journal of International Economics* 94 (2): 195–206.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124 (4): 1403–48.

Krishna, Pravin, and Devashish Mitra. 1998. "Trade Liberalization, Market Discipline and Productivity Growth: New Evidence from India." *Journal of Development Economics* 56 (2): 447–62.

Krugman, Paul R. 1979. "Increasing Returns, Monopolistic Competition, and International Trade." *Journal of International Economics* 9 (4): 469–79.

Levinsohn, James. 1993. "Testing the Imports-as-Market-Discipline Hypothesis." *Journal of International Economics* 35 (1-2): 1–22.

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

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

Ministry of Economic Affairs, Taiwan. 2000–2004. "Taiwanese Annual Survey of Manufacturing." Factory Operation Census 2000–2004, Department of Statistics. http://dmz9.moea.gov.tw/gmweb/investigate/InvestigateG.aspx?lang=E (accessed August 17, 2015).

Nelsen, Roger B. 2006. *An Introduction to Copulas.* 2nd ed. New York: Springer.

Peters, Michael. 2013. "Heterogeneous Mark-Ups, Growth and Endogenous Misallocation." http://eprints.lse.ac.uk/54254/1/Peters_heterogeneous_mark-ups_2013.pdf (accessed August 17, 2015).

Restuccia, Diego, and Richard Rogerson. 2008. "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments." *Review of Economic Dynamics* 11 (4): 707–20.

Simonovska, Ina, and Michael E. Waugh. 2014. "The Elasticity of Trade: Estimates and Evidence." *Journal of International Economics* 92 (1): 34–50.

Tybout, James R. 2003. "Plant- and Firm-Level Evidence on 'New' Trade Theories." In *Handbook of International Trade*, edited by E. Kwan Choi and James Harrigan, 388–415. Malden, MA: Blackwell Publishing Ltd.

Zhelobodko, Evgeny, Sergey Kokovin, Mathieu Parenti, and Jacques-Francois Thisse. 2012. "Monopolistic Competition: Beyond the Constant Elasticity of Substitution." *Econometrica* 80 (6): 2765–84.