Imported Inputs and Productivity[1]

László Halpern

Institute of Economics, Hungarian Academy of Sciences and CEPR

Miklós Koren

Central European University, IEHAS and CEPR

Adam Szeidl

Department of Economics, University of California - Berkeley

April 2009

[1]E-mail addresses: halpern@econ.core.hu, korenm@ceu.hu and szeidl@econ.berkeley.edu.
We thank Pol Antràs, Péter Benczúr, Christian Broda, Jan De Loecker, Gita Gopinath, Penny Goldberg, Elhanan Helpman, Marc Melitz, Ariel Pakes, Roberto Rigobon, John Romalis, David Weinstein and seminar participants for helpful comments. This research was supported by a grant of the Global Development Network (RRC IV-061) and of the Hungarian Scientific Research Fund (T048444).

Abstract

How do imported inputs affect firm productivity? We address this question by estimating a structural model of importers using product-level data for all Hungarian manufacturing firms during 1992-2003. We have three main findings. (1) Imported inputs have large productivity effects: increasing the share of imported goods from 0 to 100 percent increases productivity by 11 percent. (2) About 60% of this gain is due to imperfect substitution, i.e., the idea that combining different inputs is "more than the sum of the parts." This is consistent with Hirschman's (1958) view about the importance of complementarities along a production chain for economic development. (3) Tariff cuts have a highly non-linear effect on productivity, due to firm entry into import markets for new varieties. This non-linearity can rationalize differences between estimated tariff effects in different studies, and shows how firm level analysis helps understand macro facts. Our structural approach can also be used for counterfactual policy analysis, and to study the different implications of the quality and complementarity mechanisms.

1 Introduction

A body of work in international economics suggests that foreign trade has large positive effects on income.[1] Recent research shows that intermediate inputs may play an important role in realizing these income gains. For example, Amiti and Konings (2007) document that the productivity gain from cutting tariffs on intermediate goods are twice as big as those from comparable cuts for final goods in Indonesia. Similarly, Goldberg, Khandelwal, Pavcnik and Topalova (2008) show that access to new intermediate inputs produces substantial gains in India. More generally, Jones (2009) discusses the importance of intermediate inputs for economic development, and argues that they can help explain large income differences across countries.

How do intermediate goods affect productivity? The literature has proposed two broad channels, which we refer to as the quality and complementarity mechanisms. Quality simply means that imported inputs are better than their domestic counterparts, a mechanism frequently alluded to in the endogenous growth literature (e.g., Grossman and Helpman (1991)). Complementarity is the idea that combining different intermediate inputs create gains that are more than the "sum of the parts." These gains could come from imperfect substitution across goods, as in the love-of-variety setting of Krugman (1979) and Ethier (1982), through Kremer's (1993) O-ring theory of economic development, or through learning spillovers between foreign and domestic goods (e.g., Keller (2004)). As Jones (2009) emphasizes, in equilibrium either of these channels can create additional productivity gains through an income multiplier.

In this paper, we use Hungarian firm-level data to better understand the productivity effect of imported inputs in practice. A novelty of our approach is that we estimate a structural model of importing firms. This allows us to identify the mechanism driving productivity gains, examine other results in the literature through the lens of a model, and explore the implications of our estimates for trade policies.

Our starting point is a unique dataset that contains detailed product-level information for all imported goods for the entire universe of Hungarian firms during 1992-2003. Motivated by basic stylized facts in these data, we formulate a model of importer-producer firms who use differentiated inputs to produce a final good. In this model, intermediate goods affect output both through a quality and a complementarity channel. The model implies a firm-level production function where output depends on the usual factors as well as a term related to the number of intermediate imported goods, which reflects a combination of both mechanisms. We estimate this production function by adapting the Olley and Pakes (1996) methodology to our setting, which deals with the reverse causality and endogenous exit problems that plague productivity estimation. We find that imports generate substantial

[1] See the cross-country studies by Romer (1987), Coe and Helpman (1995), Barro (1997) and Frankel and Romer (1999).

gains: our baseline estimates imply that increasing the fraction of goods imported from 0 to 100 percent would increase productivity by 11 percent.

We then turn to the question of disentangling the quality and complementarity effects. This is where the structural nature of the model is useful. Our identification comes from noting that the strength of the complementarity mechanism governs the elasticity of import demand to quality. If foreign and domestic intermediate goods are perfect substitutes, then even small quality improvements result in large changes in import demand. Conversely, when complementarities are strong, demand responds little even to substantial quality improvements, because inputs have to be combined in "just the right proportions." This link between the elasticity of import demand and the role of complementarities is also exploited by Feenstra (1994) and Broda and Weinstein (2006) in country-level data. Our estimates imply that about 60% of firms' productivity gains from importing are due to the complementarity channel. This result provides micro-level evidence that imperfect substitution across goods has quantitatively large productivity effects. More broadly, our findings lend empirical support to an old view in development economics (Hirschman 1958) revisited by Kremer (1993) and Jones (2009), that complementarities along different parts of the production chain play an important role for output and growth.

A key advantage of our semi-structural approach is that we can use the estimated parameters to conduct counterfactual experiments. We analyze the effects of tariffs, a major tool for conducting trade policy in practice, in our estimated economy. Our main finding is that tariff changes have highly non-linear effects on productivity, due to the extensive margin through which firms enter into import markets for new varieties. Intuitively, the effect of a small tariff reduction will often be realized only through the increased foreign purchases of existing importers. In contrast, a larger reduction can prompt non-importers to start buying from abroad, and also current importers to increase the set of varieties that they import.[2]

The non-linear effect of tariff cuts also implies that they are most beneficial when the volume of firms who are at the margin of entering import markets is the largest. Given the empirical shape of the firm size distribution, this happens when a country is at an intermediate degree of openness, so that many mid-sized companies are about to become importers. This observation helps explain why different studies in the literature found different effects of input tariffs on productivity. For example, Muendler (2004) estimates small effects using data from Brazil. However, before the tariff cuts, Brazilian importers faced substantial barriers, suggesting that the productivity gains following the change were harvested by a few big firms: most companies in the economy did not become importers overnight. In contrast, Amiti and Konings (2007) and Kasahara and Rodrigue (2008) find larger effects in Indonesia and Chile, both of which were relatively open economies during the time periods studied.

[2]This large role of the extensive margin parallels the findings of Broda, Greenfield and Weinstein (2006) and Goldberg et al. (2008), who emphasize the productivity gains from new goods.

The idea that heterogenous firms respond differently to trade policies is also consistent with Konings and Vandenbussche (2008), who find differential productivity effects of antidumping protection. A broader lesson is that firm level analysis can be important to understand the aggregate implications of economic policies, because firm responses may vary depending on size and other characteristics, such as importing status.

Our tariff experiments also show that the quality and complementarity mechanisms have different implications for trade policy. In particular, the relative demand for domestic versus foreign inputs in response to a tariff cut is different under the two channels. When quality is important, a tariff reduction brings about large import substitution, hurting domestic intermediate good producers. Conversely, when complementarities matter, a tariff cut reduces the demand for domestic goods much less, because they must be combined with foreign goods to maximize output. Our baseline estimates about the importance of the complementarity channel therefore suggests that concerns by domestic producers about redistributive losses due to import competition may be misguided. Tariff cuts may in fact boost the demand for domestic intermediate goods due to increased productivity. The larger point is that identifying the mechanism through which trade policies affect the economy helps understand the consequences of those policies in other economic areas, such as domestic input production.

Besides the papers cited above, we also build on a large literature on trade and growth. Much of this work is reviewed by Hallak and Levinsohn (2008), who identify the the estimation of structural models in micro data as the next step in the research agenda. This is the approach we implement in this paper. Our work is also related to recent empirical research that studies firms' behavior in international markets, including Bernard, Jensen and Schott (2009) for the U.S. and Eaton, Kortum and Kramarz (2004) for France. Tybout (2003) summarizes earlier plant and firm level empirical work testing theories of international trade. One weakness of our analysis is that we ignore the effect of imported capital goods. Caselli and Wilson (2004) show that complementarities between the embedded technology of imported capital and domestic labor have positive productivity effects, reinforcing our finding about imperfect substitution between intermediate inputs.

The rest of this paper is organized as follows. Section 2 describes our data and collects several stylized facts about importers in Hungary. Building on these stylized facts, Section 3 develops a simple model of importer-producers. Section 4 describes the estimation procedure and results. Section 5 conducts our counterfactual experiments, Section 6 discusses some extensions, and Section 7 concludes.

2 Data

2.1 Data description

Our data consists of a panel of Hungarian manufacturing companies from 1992 to 2003. This data comes from two sources: the Hungarian Customs Statistics, and firms' balance sheets and earnings statements. The Customs Statistics data contain firms' annual exports and imports at the disaggregate level of 6-digit Harmonized System (HS) product categories (5,200 categories). We aggregate the data up to the 4-digit level (1,300 categories) because the 6-digit classification is noisy.[3] In the rest of the paper, the terms "product" and "good" refer to a HS4 category. Using the product level classification of Rauch (1999), these goods can be further broken down into homogenous and differentiated products. We will make use of this categorization in the empirical analysis.

There are 32,174 firms in our dataset. Table 1 reports summary statistics for several key variables in our sample. Column 1 reports the total number of firms, the number of observation, and the mean of several variables. Columns 2 and 3 break down these numbers by importing and non-importing firms.

Table 1: Descriptive statistics

	Total	**Non-importers**	**Importers**
Observations	174,726	111,213	63,513
Firms	32,174	26,777	14,535
Employment	52.02	17.27	112.86
Sales (thousand USD)	3,533.79	527.70	8,797.53
Capital per worker (thousand USD)	18.28	15.81	22.30
Sales per worker (thousand USD)	62.24	48.66	85.36
Material share	0.64	0.63	0.67
Exporter dummy	0.36	0.16	0.70
Export share	0.15	0.05	0.31
Importer dummy	0.36		
Import share	0.11		0.28
Number of imported products	3.86		10.63
Foreign owned	0.18	0.09	0.35
State owned	0.03	0.03	0.03

One key limitation of the data is that we do not have product-level information for domestic inputs used by the firms in our sample. As we explain below, the structural approach helps address this problem: in our model the effect of imports on productivity can be estimated even in the absence of product-level breakdown for domestic inputs, as long as such data for foreign inputs is available.

2.2 Stylized facts about firm-level imports

We now turn to document five basic facts about the importing behavior of firms in our data.

[3]Firms often switch their main export product at the 6-digit level; this happens infrequently at 4 digits.

Fact 1. There is substantial heterogeneity in the import patterns of firms: 64 percent of firms do not import at all. Importers are larger and more productive than non-importers.

As seen from Table 1, most firms do not import at all. Comparing columns 2 and 3 shows that importers employ about 6 times as many workers and sell about 12 times as much. They are also more likely to export and to be foreign owned. This is consistent with Bernard et al. (2009), who document that "globally engaged firms" in the U.S. exhibit superior performance along a number of dimensions.

Fact 2. Larger firms import more kinds of products.

The typical importer buys 11 different products from abroad (Table 1), which represents 28 percent of overall intermediate input costs. Even importers differ substantially in the intensity of imports. Figure 1 shows the relationship between importing behavior and firm size by plotting the number of imported products (HS4 categories) as a function of the log employment, for foreign and domestic firms. The lines correspond to nonparametric estimates of the relationship between number of products and employment. The number of imported products sharply increases in size: doubling firm size would increase the number of imported products by 30 percent. Foreign firms are more internationally engaged: controlling for firm size, they tend to import about twice as many products as domestic ones.

These patterns are consistent with a model where entry in import markets entails a fixed cost, perhaps because it requires establishing and maintaining business connections. Larger firms profit more from buying a given product and hence find it easier to overcome the fixed cost. Foreign firms who are more likely to have a business network abroad plausibly face lower fixed costs, which can explain why they import more varieties.

Fact 3. Imports are concentrated on a few products, and firms spend little on their remaining imports.

The expenditure on additional imported product varieties declines sharply with the number of goods imported: for example, the average firm spends 45% of its import budget on the largest product category and only 4% on the fifth largest category. This finding implies that we cannot treat product categories symmetrically in our model: we need to take into account diminishing demand for additional imported varieties.[4]

Fact 4. Over time, imports have mostly increased on the "extensive margin": more firms started importing.

Table 2 shows how the fraction of importing firms as well as the share of imports in intermediate expenditure have changed over the sample period. The fraction of importers increased from 36 percent in 1992 to 44 percent in 2003. Conditional on importing, the share of imported intermediates displays no clear time trend.

[4]See Hummels and Lugovskyy (2004) for a model where the marginal utility of additional varieties declines.

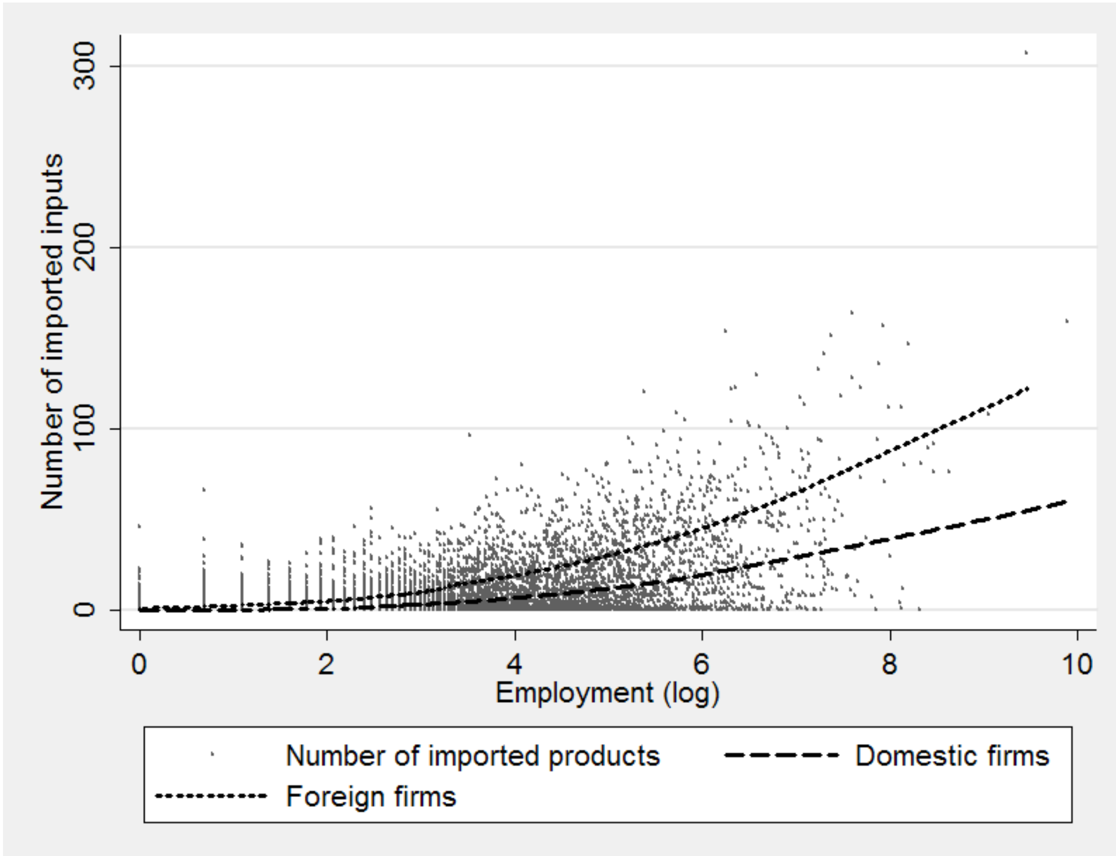

Figure 1: Number of imported inputs and firm size

Fact 5. Growing firms enter new import markets and shrinking firms exit their existing import markets.

Figure 2 plots the proportional change in the number of products imported against sales growth. To allow for zeros, product growth is defined as $\Delta N_{it}/(N_{it} + N_{i,t-1})$. A value of -1 means that the firm has stopped importing, ($N_{it} = 0$), whereas $+1$ means that the firm is a new importer ($N_{i,t-1} = 0$). The key point of the figure is that there is no visible asymmetry between growing and shrinking firms. This finding suggests the costs of participating in import markets are best modeled as fixed costs and not sunk costs: firms who want to continue importing a particular good need to pay the cost every period.

3 A Model of Intermediate Goods and Productivity

Motivated by these stylized facts, we now develop a model where firms can import or purchase domestically differentiated intermediate goods.

Table 2: Trends in import patterns

Year	Fraction of importers	Import as a share of intermediates
1992	0.36	0.29
1993	0.35	0.30
1994	0.34	0.31
1995	0.33	0.32
1996	0.32	0.34
1997	0.31	0.34
1998	0.31	0.34
1999	0.31	0.32
2000	0.42	0.30
2001	0.44	0.27
2002	0.44	0.25
2003	0.44	0.25

3.1 Production technology

Firms use capital, labor and materials in their production process. Total output is determined by the production function

$$Y = \Omega K^\alpha L^\beta \prod_{i=1}^{\mathcal{N}} X_i^{\gamma_i}, \tag{1}$$

where K denotes capital inputs, L labor inputs, X_i is the amount of composite good i and Ω is Hicks neutral total factor productivity (TFP). Here γ_i is the Cobb-Douglas weight of composite good i in the production process. The total weight of all intermediate goods is $\gamma = \sum_i \gamma_i$. We also assume that each good X_i is assembled from a combination of two varieties, a foreign and a domestic one:

$$X_i = \left[(B_i X_{iF})^{\frac{\theta-1}{\theta}} + X_{iH}^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}, \tag{2}$$

where X_{iF} and X_{iH} are the quantity of foreign and domestic inputs, and θ is the elasticity of substitution.

This specification for production incorporates both the quality and complementarity channels discussed in the introduction. $B_i > 0$ measures the quality of the foreign good i relative to its domestic counterpart. Given that Hungary is a developing country with a majority of imports coming from advanced economies, we expect foreign goods to enjoy a quality

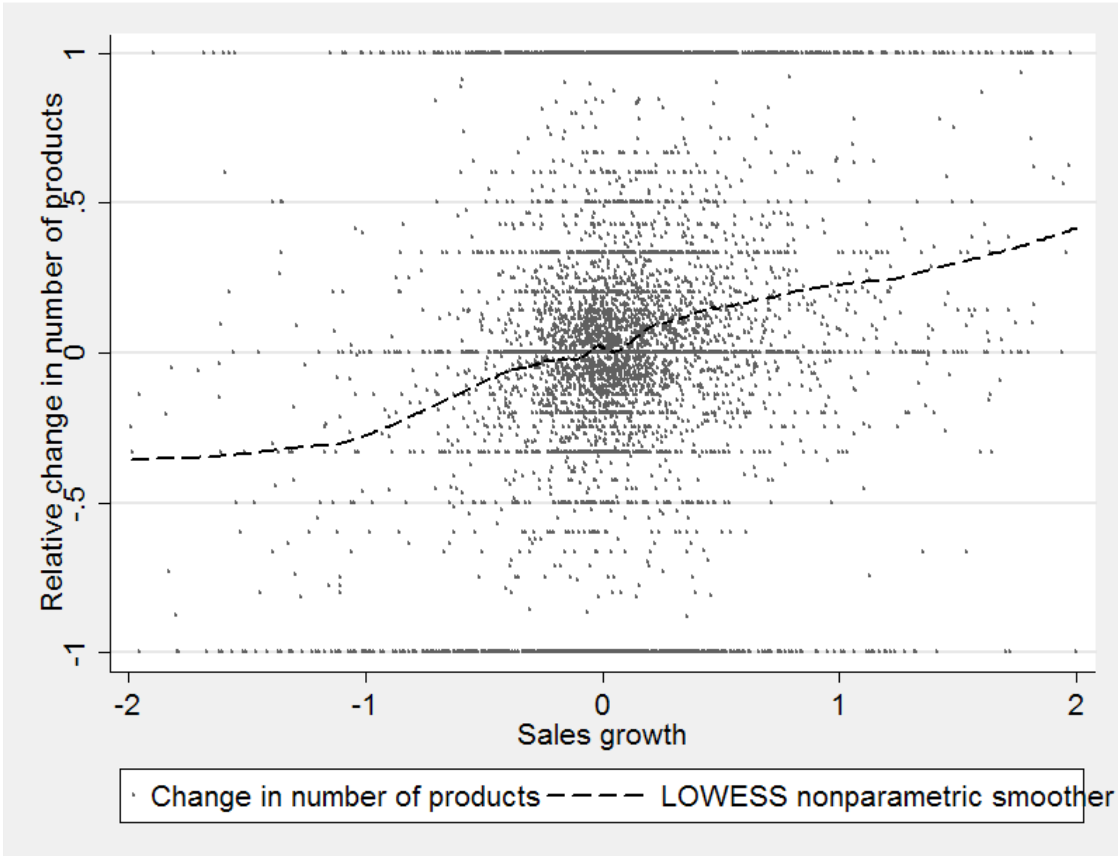

Figure 2: Product market entry and firm growth

advantage. Complementarity is measured by θ, the elasticity of substitution between foreign goods and their domestic counterparts. The higher θ, the more substitutable these inputs are, reducing the gains from combining them. Given the Cobb-Douglas aggregation across goods, a higher θ also means greater substitution between different imported inputs; in the limit, when θ grows without bound, different foreign goods also become perfect substitutes.

A key advantage of this specification is that it retains the tractability of the Cobb-Douglas model, while allowing for flexibility in the degree of substitution both between foreign and domestic varieties and between different imported goods. This flexibility is central to our analysis, because we want to estimate the importance of complementarities from the data. A weakness is that a single parameter θ governs both within and across-good substitution; but this must be weighted against the analytical convenience of the Cobb-Douglas model, which is the standard theoretical framework in production function estimation. Like in the production function employed by Jones (2009) and Kremer (1993), our Cobb-Douglas specification implies that all inputs are essential: if the use of any factor or intermediate good falls to zero, no output is produced.

3.2 The productivity gains from importing

The standard approach to measuring productivity is to hold fixed all factors and attribute any additional variation in output to productivity. In our model, this approach runs into the difficulty that the set and composition of intermediate inputs varies across firms depending on their access to import markets. To deal with this problem, we extract a measure of productivity from the marginal cost schedule of the firm. Intuitively, firms with higher productivity should produce an additional unit of the final good at a lower marginal cost given input prices.

The marginal cost function of the firm, ignoring constants, is

$$C(Y, R, W, \{P_i\}) = Y^{1/(\alpha+\beta+\gamma)-1}\Omega^{-1/(\alpha+\beta+\gamma)}R^{\alpha/(\alpha+\beta+\gamma)}W^{\beta/(\alpha+\beta+\gamma)}\prod_{i=1}^{\mathcal{N}} P_i^{\gamma_i/(\alpha+\beta+\gamma)} \tag{3}$$

where R is the cost of capital, W is the wage rate, and P_i is the cost-minimizing price of the bundle X_i conditional on the entry decision in import markets. Changing units so that the domestic price of all goods is normalized to $p_{iH} = 1$, the effective price of good i is $P_i = 1$ if the firm only uses domestic inputs, but $P_i < 1$ if the firm also imports, because combining the two inputs results in the same composite X_i at a lower cost. Formally,

$$P_i = \begin{cases} \left[p_{iH} + (p_{iF}/B_i)^{1-\theta}\right]^{1/(1-\theta)} & \text{if } i \text{ is imported,} \\ p_{iH} = 1 & \text{otherwise.} \end{cases} \tag{4}$$

We denote by a_i the percentage decrease in the cost of bundle i if it uses imported products,

$$a_i = \frac{\ln\left[1 + (B_i/p_{iF})^{\theta-1}\right]}{\theta - 1}. \tag{5}$$

This expression summarizes both the quality and complementarity channels. The percentage gain from importing good i increases in its quality advantage, B_i, and decreases in the degree of substitution, θ. Figure 3 illustrates the productivity gains from importing good i as a function of the import share. In the left of the figure, when this share is zero (no imports), the associated gain is zero. On the right, when the import share is one, TFP gains are positive, reflecting the quality advantage of foreign goods. However, the highest productivity gains are realized at an intermediate import share: given imperfect substitution, inputs must to be combined in the right proportion to maximize output.

Let χ_i be an indicator variable that equals one if product i is imported and zero otherwise. Then the log of the cost function (3) can be written as

$$\ln C = \frac{1}{\alpha+\beta+\gamma}\left[(1-\alpha-\beta-\gamma)y + \alpha r + \beta w - -\omega - \sum_{i:\chi_i=1}\gamma_i a_i\right].$$

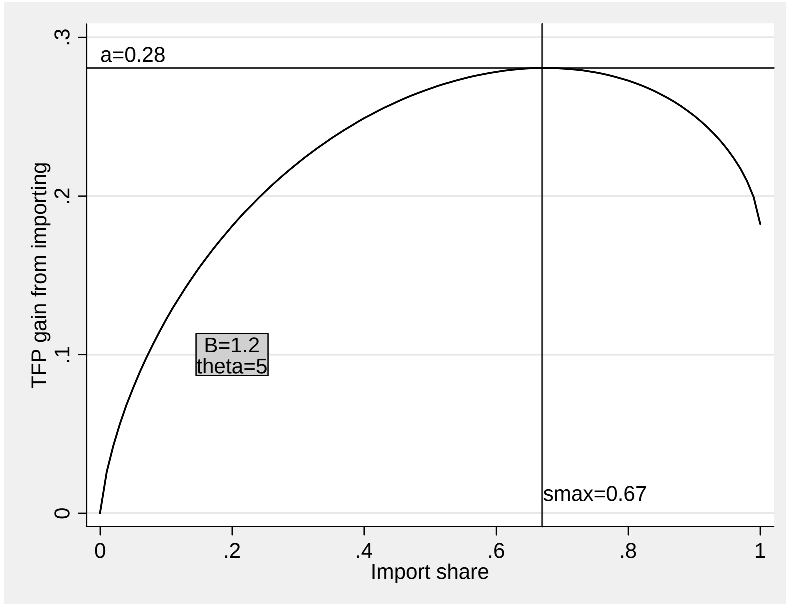

Figure 3: Productivity gains from importing

This equation indicates that holding fixed factor prices, all additional variation in costs is driven by the term

$$\omega + \sum_{i:\chi_i=1} \gamma_i a_i. \tag{6}$$

We interpret this quantity as total factor productivity, where the efficiency gains due to imports are captured by $\sum_{i:\chi_i=1} \gamma_i a_i$, while ω represents "residual" productivity. Intuitively, the productivity gain from importing variety i is proportional to the cost saving from imports, a_i, as well as the Cobb-Douglas weight of product i in the total output, γ_i.

The TFP measure (6) can also be used to express output in a conventional production function. To see how, define $M = \sum_i P_i X_i$ as the total dollar spending on intermediate inputs. Given our Cobb-Douglas production function, M must equal the price index of all inputs times the Cobb-Douglas aggregate of the inputs:

$$M = \prod_{i=1}^{\mathcal{N}} P_i^{\gamma_i/\gamma} \prod_{i=1}^{\mathcal{N}} X_i^{\gamma_i/\gamma}.$$

Taking logs in this expression, substituting the definition of the cost-minimizing bundles and combining it with production function (1) yields

$$y = \alpha k + \beta l + \gamma m + a \sum_{i:\chi_i=1} \gamma_i + \omega. \tag{7}$$

This equation has the form as a standard production function, where log output depends on various factors. The first three terms on the right hand side measure the contribution of capital, labor and intermediate inputs to total output (we denote $m = \ln M$). The novelty lies in the last two terms, which, by equation (6) decompose total factor productivity into the effect of imports and residual productivity.

3.3 The decision to import

Equation (7) shows that participating in import markets has a positive effect on output. Absent other considerations, firms would thus find it optimal to import a positive amount of all foreign input varieties. But it is clear from the data that most firms only import only a few product varieties and some firms do not import at all. One plausible reason for this empirical pattern is that, similarly to Melitz (2003) for exports, importing any given variety i has an associated fixed cost. For example, importing a product might require maintaining a business connection, which has a cost even if the purchased quantity is small. These fixed costs are likely to be recurring rather than sunk because, by Fact 5, there is substantial variation both upward and downward in the set of products a firm imports over time.

Motivated by these observations, we assume that there is an initial fixed cost F that a firm must pay in order to access foreign markets; and that entering the import market for each variety has an additional fixed cost f. We assume that f is constant across products within a firm, but potentially different across firms: some companies, including those owned by foreigners, might find it cheaper to enter import markets. These costs are due every period when the firm is importing.

Conditional on importing, the dollar cost reduction from importing product i is $a_i \gamma_i Y$, and hence an importer firm will choose to import i if and only if these gains exceed the fixed cost

$$\gamma_i a_i Y \geq f. \tag{8}$$

To characterize the decision to enter import markets, order products so that $\gamma_1 a_1 \geq \gamma_2 a_2 \geq$ ..., and let n^* denote the final index where (8) still holds, i.e., the "last" product that the firm chooses to import: $n^* = n^*(Y, f) = \max\{n : \gamma_n a_n Y \geq f\}$. Clearly, firms who are larger or have lower fixed costs import more product varieties.

To characterize the decision to enter import markets, we denote the fraction of potential productivity gains realized through importing goods $1, ..., n$ by

$$G_n = \frac{\sum_{i=1}^{n} \gamma_i a_i}{\sum_{i=1}^{N} \gamma_i a_i} = \frac{\sum_{i=1}^{n} \gamma_i a_i}{\gamma \bar{a}} \tag{9}$$

where $\bar{a} = \sum_{i=1}^{N} \gamma_i a_i / \gamma$ is the average potential cost-saving from importing. It follows that a firm with size Y, fixed costs f and F will choose to enter import markets if and only if

$$\gamma \bar{a} G_{n^*} > n^* f + F, \tag{10}$$

that is, if the total gains from importing exceed the fixed costs of entering import markets and importing the optimal number of varieties. The actual number of imported varieties n

is then equal to n^* if (10) holds, and is zero otherwise. In particular, large firms are more likely to import (fact 1), and, conditional on importing, they import more kinds of products (fact 2). The same holds for firms with low fixed costs.

The endogenous import decision is reflected both in firm output and import demand. Using the notation just introduced, the production function can be re-written as

$$y = \alpha k + \beta l + \gamma m + \gamma \bar{a} G_n + \omega \tag{11}$$

where G_n measures the realized productivity gains from importing. Total spending on imports also depends on the set of imported varieties:

$$\frac{X_F}{Y} = \frac{\sum_{i=1}^{\mathcal{N}} p_{iF} X_{iF}}{Y} = \sum_{i=1}^{n} \gamma_i s_i, \tag{12}$$

where $s_i = (B_i/p_{iF})^{\theta-1}/(1 + (B_i/p_{iF})^{\theta-1})$ is the optimal share of imports in product i. Introducing the notation $\bar{s} = \sum_{i=1}^{N} \gamma_i s_i/\gamma$ for the average import share across goods and

$$H_n = \sum_{i=1}^{n} \gamma_i s_i / \gamma \bar{s},$$

for the value-weighted fraction of products that are imported, import demand can be written as

$$\frac{X_F}{Y} = \gamma \bar{s} \cdot \frac{\sum_{i=1}^{n} \gamma_i s_i}{\gamma \bar{s}} = \gamma \bar{s} \cdot H_n. \tag{13}$$

Equations (8) - (13) completely characterize firm behavior in our model.

4 Estimation

4.1 Estimating equations

Our goal is to estimate the structural parameters of the production function α, β, $\{\gamma_i\}$, B_i and θ, as well as the firm-specific fixed costs F and f, which govern the entry into import markets. Our first estimating equation is the production function, which we write as

$$y = \alpha k + \beta l + \gamma m + \delta G_n + \tilde{\omega}, \tag{14}$$

where $\delta = \gamma \bar{a}$ and the error term $\tilde{\omega}$ is residual productivity, part of which may be observable by the firm before making input choices.

The second estimating equation is the firm's import demand. We first note that up to a first order approximation around $a_i = 0$, $s_i \approx (\theta - 1)a_i$, so that $H_n \approx G_n$.[5] We can then use equation (13) to write

$$\frac{\sum_{i=1}^{\mathcal{N}} p_{iF} X_{iF}}{Y} \approx \gamma \bar{s} G_n + \tilde{u} \tag{15}$$

[5]This approximation holds for small a_i *or* if a_i varies little across products. The equation $H_n = G_n$ holds *exactly* when $a_i = a$ for all i.

where the error term $\tilde{u}$ captures measurement error and other unmodelled variation in imports. Firms that import a greater number of products spend a larger fraction of their output on imported inputs. Our third equation is an inequality characterizing endogenous entry to import markets:

$$a_n \gamma_n Y \leq \tilde{f} < a_{n+1} \gamma_{n+1} Y, \quad (16)$$

where $\tilde{f}$ is the firm-specific per-product cost of entry in import markets, which can vary across firms and over time. Our final step in the estimation uses (10) to match the share of importing firms in the data.

For identification, we need to make some assumptions about the distribution of error terms. Estimating (14) is difficult, because $\tilde{\omega}$ is likely to be correlated with all variables on the right hand side, for example due to reverse causality problems. We deal with this endogeneity problem using the strategy of Olley and Pakes (1996). The identification assumptions required are that (i) $\tilde{\omega}$ can be written as a sum of two terms, a deterministic, monotonic function of observables such as investment, and an innovation orthogonal to other variables on the right hand side of (14); and (ii) that the first term in this decomposition of $\tilde{\omega}$ follows a first-order Markov process.[6]

To estimate equation (15), we assume that $\tilde{u}$ is classical measurement error, orthogonal to the decision to enter import markets and hence the number of imported inputs n. This allows us to use an ordinary least squares regression. One important case where this assumption is violated is when there is cross-firm heterogeneity in s, which would result in a nonzero correlation between $\tilde{u}$ and G_n. Section 6 discusses the robustness of our estimates to such heterogeneity. Finally, our estimating procedure requires no assumptions about the error term $\tilde{f}$ in equation (16), but assumes that F is constant across firms.

Given these assumptions, we proceed through the following steps. (1) We back out the "potential productivity gains function" G_n from observed import shares. (2) We estimate the production function (14). (3) We estimate the import demand to obtain $\bar{s}$, and use it to back out the quality and complementarity channels. (4) We use equation (16) and the share of importers to measure fixed costs. We now describe these steps in detail.

4.2 Estimating the potential gains function G_n

We estimate the gains function G_n from import demand. Taking expectations in (15) conditional on the number of imported products n yields

$$E\left(\frac{X_{Fjt}}{Y_{jt}} | n_{jt}\right) = \gamma \bar{s} G_n.$$

We use this equation to recover the shape of the G_n function. The conceptual idea is to average the import share across all firms who import exactly n different products in our

[6]See Olley and Pakes (1996) for a discussion of the validity of these assumptions.

data. The representativeness of our data is important for this step, otherwise we would have to worry about selection bias. We can also identify the term $\gamma\bar{s}$, because we know that G_n attains a value of 1 at the maximal number of products, $\mathcal{N}$. The parameters γ and $\bar{s}$ are not yet separately identified.

To implement the "averaging" described above, we use a non-parametric estimator of import demand as a function of n. Our benchmark specification estimates X/Y as a third-order polynomial of $\ln(1+n)$, but a lowess specification yielded similar results. The predicted value of this estimator for $n = N$ gives the estimate of $\gamma\bar{s}$. As a byproduct of these calculations, we also obtain estimates of

$$\frac{\gamma_i a_i}{\gamma \bar{a}} = G_i - G_{i-1}.$$

4.3 Estimating the production function

We estimate (14) using the Olley and Pakes (1996) approach. Due to the additional right hand side variable G_n, we need to make small changes to their procedure; Appendix A explains in detail what we do. This step yields estimates of the parameters α, β, γ, δ, and hence $\bar{a} = \delta/\gamma$ as well.

4.4 Identifying structural parameters

Given our estimates of G_n from the previous subsection, we can just estimate (15) in ordinary least squares to obtain $\bar{s}$. We now turn to explain how these estimates allow us to back out the deep parameters B and θ, which govern the strength of the quality and complementarity channels.

The model implies the following expressions for B_i and θ:

$$\ln B_i = a_i \left[1 - \frac{\ln \ s_i}{\ln(1 - \ s_i)}\right], \tag{17}$$

$$\frac{1}{\theta - 1} = \frac{a_i}{-\ln(1 - s_i)}. \tag{18}$$

The term $\ln B$ in the first equation measures the quality channel, while $1/(\theta - 1)$ is inversely related to θ and hence is increasing in the importance of complementarities. To interpret these equations, first note that both the quality and the complementarity effects are proportional to a, the total benefits from importing. This is because both channels contribute positively to the overall productivity gain from imports.

To assess the relative importance of the two channels, note that for a fixed a, the quality effect is increasing while the complementarity effect is decreasing in s. Intuitively, a higher quality differential induces firms to spend more on imports relative to domestic inputs, leading to a higher s. In contrast, higher complementarity calls for keeping imports and domestic inputs in proportion to one another, resulting in a lower s.[7] For example, in the limit when foreign and domestic inputs are no longer substitutes ($\theta = 1$), the firm would spend equal amounts on both ($s = 1/2$), regardless of the quality differential.

To summarize, a high value of a results in higher estimates for both the quality and variety channels, while the relative contribution of these channels is governed by the magnitude of s. We use this logic to identify the key parameters of the model and the contribution of the quality and variety channels to the total gains from importing.[8]

Implementing this approach runs into the problem that so far we only estimated the average values $\bar{a}$ and $\bar{s}$. In the baseline estimates, we deal with this by assuming that the quality advantage of all products is the same: $B_i = B$ for all i. This assumption implies that $a_i = \bar{a}$ and $s_i = \bar{s}$ for all products, and inverting the above equations is then straightforward. In Section 6 we discuss how the results change when we allow heterogeneity across products.

4.5 Computing fixed costs

We now turn to estimate the firm-specific fixed entry costs. Given our estimate for the productivity improvement a, data on the number of varieties imported by each firm allows us to back out an estimate f for importers. While in general (16) only allows us to identify a range in which the fixed cost must lie, for n large, i.e., firms that import many varieties, these intervals are tight and to a close approximation, we obtain $f = \gamma_n a_n Y$. Our results are less precise for small values of n.

To estimate the fixed cost of accessing foreign markets F, we use information on the share of non-importers. We first show that the realized values of f are well-approximated by a lognormal distribution conditional on firm size. Estimating parameters of this lognormal, and combining it with our estimates of the parameters in the theory, for each possible value of F, the model predicts the probability that a given sized firm will choose to enter import markets. Integrating these probabilities through the empirical firm size distribution gives us a prediction for the share of importers in the model. We set the value of F to match the empirical share of importers in this procedure.

4.6 Standard errors

We obtain standard errors for all our estimates from a bootstrap. We bootstrap the estimates by sampling firms in the data with replacement, keeping the entire time series of a firm together. We do this because the estimation of the capital coefficient relies of past and

[7] Note that as long as $B \geq 1$, s is never smaller than 1/2, because conditional on importing a given product, the share of imports must be at least 50 percent if foreign goods are not worse than domestic goods.

[8] Our model is also consistent with the Feenstra (1994) and Broda and Weinstein (2006) measures of gains from variety. They express productivity (welfare) gains as $\lambda^{1/(1-\theta)}$, where $\lambda < 1$ is the new expenditure share of old varieties. In our model, $\lambda_i = 1 - s_i$ and the productivity gains are $\exp(a_i) = (1 - s_i)^{1/(1-\theta)}$.

future information about the firm because of the dynamics of productivity, the endogeneity of exit and the use of lagged variables as instruments, and hence we cannot treat time periods as independent observations.

The bootstrap procedure is as follows. First we create 200 random samples from the data. Then, for each new sample, we repeat the estimation steps described above to fan out the empirical distribution of all parameters. The standard errors our tables report are calculated as the standard deviation of parameter estimates across replications.

4.7 Estimation results

Table 3 summarizes the results from our estimation. Column 1 shows an OLS estimate of the production function (14), while columns 2-6 report the results of the GMM procedure outlined above, in various specifications.[9] Note that in all specifications, the dependent variable is total sales, not value added; hence the large coefficients of material costs and the relatively small coefficients of capital and labor. Also, all columns include year and industry fixed effects; thus we are identifying from within-year and within-industry variation across firms.

We expect the OLS coefficients of freely adjustable inputs, like labor, material, and imported inputs, to be upward biased due to reverse causality, while the coefficient of capital is likely to be downward biased because of exit. Comparing column 1 with the other columns shows that the OLS coefficients of labor and materials are generally similar to the Olley-Pakes estimates. That is, we find that the reverse causality bias is modest in practice.

Turning to the key δ coefficient of the imported inputs variable G_n, column 1 reports a significant positive estimate: a 10 percentage point increase in G_n is associated with a 1.7 percent increase in output with the same amount of inputs. However, this relationship cannot be interpreted as a causal effect because of the endogeneity problems of the OLS regression mentioned earlier.

Column 2 is a minimal specification of our GMM procedure where we only include the terms in the production function (14) and year and industry effects. The point estimate of δ is now of 0.177: a 10 percentage point increase in the Cobb-Douglas weight of goods imported leads to a 1.8 percent increase in productivity. This estimate implies that the proportional benefit of importing a given product is $a = \delta/\gamma = 0.229$. Thus, combining foreign domestic goods of a given variety is on average 23 percent more productive than spending the same amount of money only on the domestic good.

The table also reports our estimates of the structural parameters B and θ which are backed out from our estimates for the import share s and the TFP gain a. In column 2, the

[9] The control function $g(\cdot)$ used to take out variations in unobserved productivity in the first stage of the Olley-Pakes procedure contains third-order polynomials of investment and capital, estimated separately for each year. We use capital, lagged capital and lagged imported capital goods as instruments in the second stage to identify the coefficient of k.

Table 3: Production function estimates

	(1) OLS	(2) OP	(3) Partial controls	(4) Full controls	(5) Foreign firms	(6) Domestic firms
Capital	0.018	0.016	0.020	0.021	-0.020	0.026
	(0.002)	(0.008)	(0.008)	(0.003)	(0.012)	(0.008)
Labor	0.198	0.201	0.216	0.209	0.262	0.201
	(0.003)	(0.003)	(0.017)	(0.003)	(0.007)	(0.017)
Materials	0.780	0.773	0.769	0.777	0.771	0.774
	(0.003)	(0.003)	(0.158)	(0.003)	(0.010)	(0.158)
Imported inputs	0.169	0.177	0.142	0.107	0.132	0.148
	(0.006)	(0.007)	(0.017)	(0.011)	(0.018)	(0.017)
TFP gain (*a*)	0.217	0.229	0.185	0.136	0.171	0.192
	(0.008)	(0.010)	(0.022)	(0.014)	(0.024)	(0.022)
Optimal share (*s*)	0.587	0.587	0.587	0.587	0.553	0.924
	(0.016)	(0.016)	(0.016)	(0.016)	(0.011)	(0.020)
Relative quality (B)	1.098	1.095	1.076	1.056	1.046	1.204
	(0.020)	(0.039)	(0.051)	(0.014)	(0.011)	(0.051)
Elasticity (theta)	5.07	4.86	5.79	7.48	5.72	14.49
	(0.14)	(0.07)	(3.05)	(0.64)	(0.75)	(3.05)
Year/industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Other controls:						
Foreign ownership			Yes	Yes	Yes	Yes
Export participation				Yes	Yes	Yes
Number of observations	159,737	159,737	159,737	159,737	159,737	159,737
Number of firms	32,173	32,173	32,173	32,173	32,173	32,173

Notes: Table reports the parameters estimated by ordinary least squares (first column) and the Olley and Pakes (1996) procedure (all remaining columns). Dependent variable is log output. Bootstrapped standard errors (in parantheses) are clustered by firm. In the OP estimation we proxy productivity by investment, controlling for capital, industries and years. Instruments include capital, lagged capital, and lagged imported capital.

quality differential $B = 1.095$ is significantly greater than one. Thus foreign goods are about 9.5 percent better quality than their domestic counterparts. Given that the total gain from importing a particular variety is 22.9 percent, the quality channel is responsible for about 40% of the overall gain in productivity due to importing a foreign good. The remaining 60% of the total gain comes from the imperfect substitution of the domestic and foreign inputs. Consistent with this, we find that the elasticity of substitution between domestic and foreign goods is relatively small at $\theta = 4.9$.

Columns 3-6 verify the robustness of these basic results in different specifications. In column 3 we control for foreign ownership. We have seen in Figure 1 that foreign firms import more kinds of products, and hence have higher G_n. If they are also more productive for other reasons, such as better management or access to capital, then our estimates of δ may be biased.[10] Column 3 shows a productivity gain from imported inputs of 0.142, which is still highly significant both statistically and economically. The share of productivity gains coming from the complementarity channel is still around 60%.

Column 4 is our preferred specification, which also controls for the export market participation of the firm by including the export share in total revenue in the first-stage regression. Because exporting and importing are correlated across firms, it is important to identify the effects of imports separately from exports. The gains from importing in this specification are 0.107, which implies a a significant quality advantage of $B = 1.056$ and an elasticity of substitution of 7.5. Again, about 60% of the total gain comes from imperfect substitution. Finally, columns 5 and 6 estimate the model separately for foreign and domestic firms. These specifications also show that both types of firms gain significantly from importing; we find that foreign firms gain more on the complementarity channel, while domestic firms enjoy higher quality gains.

A robust finding of these regressions is that imperfect substitution plays a large role in the productivity effect of imports. This result is consistent with an old view in economic development about the role of weak links along the production chain, which has recently been revisited by several papers. Jones (2009) explains the logic as follows: "high productivity in a firm requires a high level of performance along a large number of dimensions. Textile producers require raw materials, knitting machines, a healthy and trained labor force, knowledge of how to produce, security, business licenses, transportation networks, electricity, etc. These inputs enter in a complementary fashion, in the sense that problems with any input can substantially reduce overall output. Without electricity or production knowledge or raw materials or security or business licenses, production is likely to be severely curtailed." Our findings provide direct evidence for the importance of this sort of imperfect substitution in the context of combining different intermediate inputs. These results results suggest that further analysis of the complementarity mechanism, focusing both on complementarities be-

[10]We implement this by letting the control function $g(\cdot)$ include third-order polynomials of investment and capital, estimated separately for ownership type for each year.

tween different production stages as well as between different goods, may a useful area of research for understanding economic development.

Fixed costs. Our GMM approach also yields estimates of the fixed cost of importing, which are reported in Table 4. The fixed cost of entering import markets (F) is around \$13,500 (at 2000 prices). Recall, this number is estimated by asking the model to match the high fraction of non-importers in the data. The average per-product fixed costs (f) are \$3,300. This estimate is consistent with the large variation in the number of products firms import: non-negligible product-level costs are required to prevent firms from importing the maximum number of products. There appears to be substantial cross-firm variation in f, as the median value is only \$1,000.

Table 4: Fixed costs of importing

Fixed cost	**(2000 USD)**
Fixed cost of importing	13,500
Per-product cost	
median	1,000
mean	3,300

Notes: Table shows the fixed costs of importing in 2000 dollars, estimated as described in Section 4.5.

5 Counterfactual Experiments

We now turn to use our estimated model to predict the effects of tariff changes. We consider a policy where all imported products are subjected to a uniform tariff change of size τ, which effectively reduces the price of foreign inputs by a factor of $(1 + \tau)$. For example, $\tau = 0.1$ means a tariff cut of ten percentage points. What is the effect of this change on firm productivity? We address this question in a partial equilibrium setup, where the tariff change reduces the price of all foreign inputs, but does not affect the prices of other inputs and the final good. These assumptions can be justified if the country is a small open economy, and tariffs only affect a small industry. In this case, the prices of foreign inputs and output are determined on international markets, and due to the small size of the industry domestic general equilibrium effects can be safely ignored.

Intensive margin. One effect of a tariff cut is that importers increase spending on foreign products they are currently purchasing, since their effective prices have fallen. It can be shown that as a result of this change, the contribution to productivity of each good increases to

$$a' = \frac{1}{\theta - 1} \ln \left[1 + (1+\tau)^{\theta-1} B^{\theta-1} \right] > a.$$

Note that if τ is small, the increase in a good's contribution to productivity is $a' - a \approx \tau s$, i.e., total gains approximately equal the price reduction on imports τ multiplied by the import share s.

Extensive margin. A second effect of a tariff cut is that firms may choose to enter more import markets. This takes place on two margins: former non-importers may choose to become importers; and former importers may expand the set of imported varieties. To compute these effects, we combine the empirical firm size distribution with our estimate of the entry cost to become an importer F and the lognormal approximation for the distribution of fixed costs per variety f. Using these estimates, for each firm size, we use (10) to compute, given new import prices, what share of firms (i.e., what share of f realizations) find it beneficial to enter import markets, and if so, for how many varieties.

We combine these calculations to evaluate the effect of tariff changes on productivity and importing activity. To highlight the different implications of the quality and complementarity mechanisms, we consider the following three scenarios. (1) The benefits of importing are entirely caused by the quality advantage of foreign goods. (2) These benefits are entirely due to imperfect substitution. (3) About 60% of the gains are caused by imperfect substitution, as in our baseline estimate. Comparing these scenarios will be useful to understand both the different implications of these mechanisms, and why it is important to estimate the true mechanism.

Results. Our findings are summarized in Figures 4-6. In these figures, the horizontal axis measures the magnitude of the tariff change: we consider both tariff increases and reductions. Figure 4 shows the effect of tariffs on average productivity in our economy, where the three lines correspond to the three scenarios explained above. The key lesson here is that curves are convex: small tariff cuts have only marginal effects on productivity, but the magnitude of the effect increases with the size of the tariff cut. In our baseline estimate, a 5 percent tariff cut results in a productivity gain of only 2.5 percent, but a tariff cut ten times as big brings about gains of 34.6 percent, which are almost fourteen times as high.

Figure 4 also shows that the effect of a given tariff cut of, say 10 percent, will depend on the initial degree of openness in the economy. When initial trade barriers are high, so that the starting point is in the right of the figure, the tariff cut has small effects, as the gains are concentrated on a few large firms. Conversely, when the economy is very open, so that most firms, measured by sales, import some of their main inputs, the benefit of the tariff cut is mostly concentrated on the intensive margin. In the intermediate range, however, the scope for productivity gains is large, because the extensive margin for new importers and new imported varieties plays a major role. Figure 5 highlights this non-linearity on the extensive margin by plotting the share of firms that are importers as a function of the tariff.

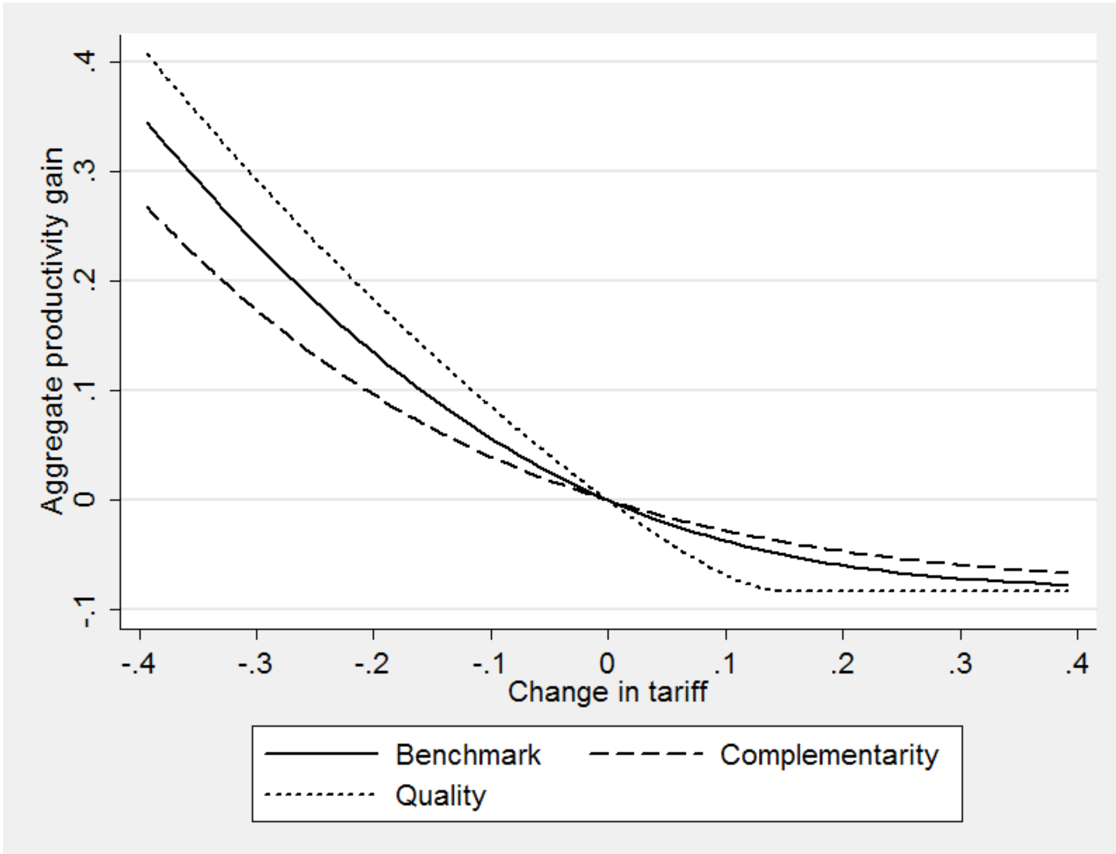

Figure 4: Productivity gains from tariff changes

The resulting inverted S-shaped curve shows the greatest sensitivity to tariff changes in the intermediate range where most firms are at the margin of becoming importers.

These non-linearities can help explain why different studies find different effects of tariff cuts on productivity. The small estimates of Muendler (2004) for Brazil may be a consequence of the fact that Brazil was a relatively closed economy prior to the tariff cuts, suggesting that, at least initially, the benefit of tariff cuts was mostly concentrated on the largest firms. Conversely, the larger effects of Amiti and Konings (2007) and Kasahara and Rodrigue (2008) are for Indonesia and Chile, both of which are fairly open economies. A broader lesson here is that understanding the aggregate effects of policies may require firm level analysis, if different firms respond differently depending on their size and on their initial importing status. This point is also reflected in the findings of Konings and Vandenbussche (2008), who show that the effects of antidumping protection differ across firms.

Comparing the quality and complementarity mechanisms. Figure 6 plots the percentage change in domestic input purchases in response to tariff changes. The key point to note in this figure is that domestic input demand is much more sensitive to tariffs when the benefit of imports comes from quality, then when it comes from complementarities. This is

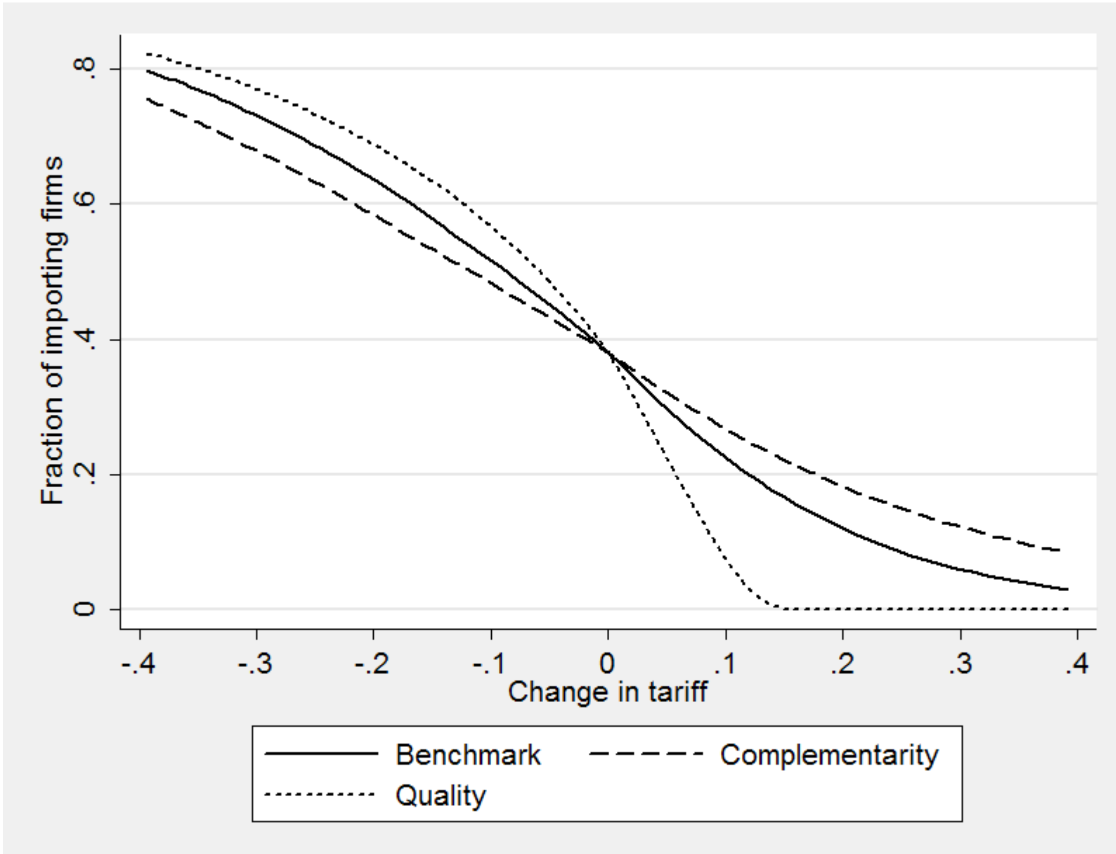

Figure 5: Tariff changes and the fraction of importers

intuitive: when foreign goods are superior to domestic, and are close to perfect substitutes, even small price changes can bring about large import substitution. In fact, this observation forms the basis of our identification of the two channels in the data. What is useful to note here is that, given our estimates, the import substitution effect is less powerful in practice, because of complementarities. One implication of this fact is that concerns about redistributive losses due to import substitution may be misguided: the demand loss is small due to complementarities in the production chain. A broader lesson is that identifying the specific mechanism driving the effect of trade policies is useful in that it helps evaluate the impact of these policies in other dimensions such as domestic input production.

6 Robustness and Extensions

We now turn to explore the robustness of the estimates to heterogeneity in the key parameters of the model and related endogeneity concerns.

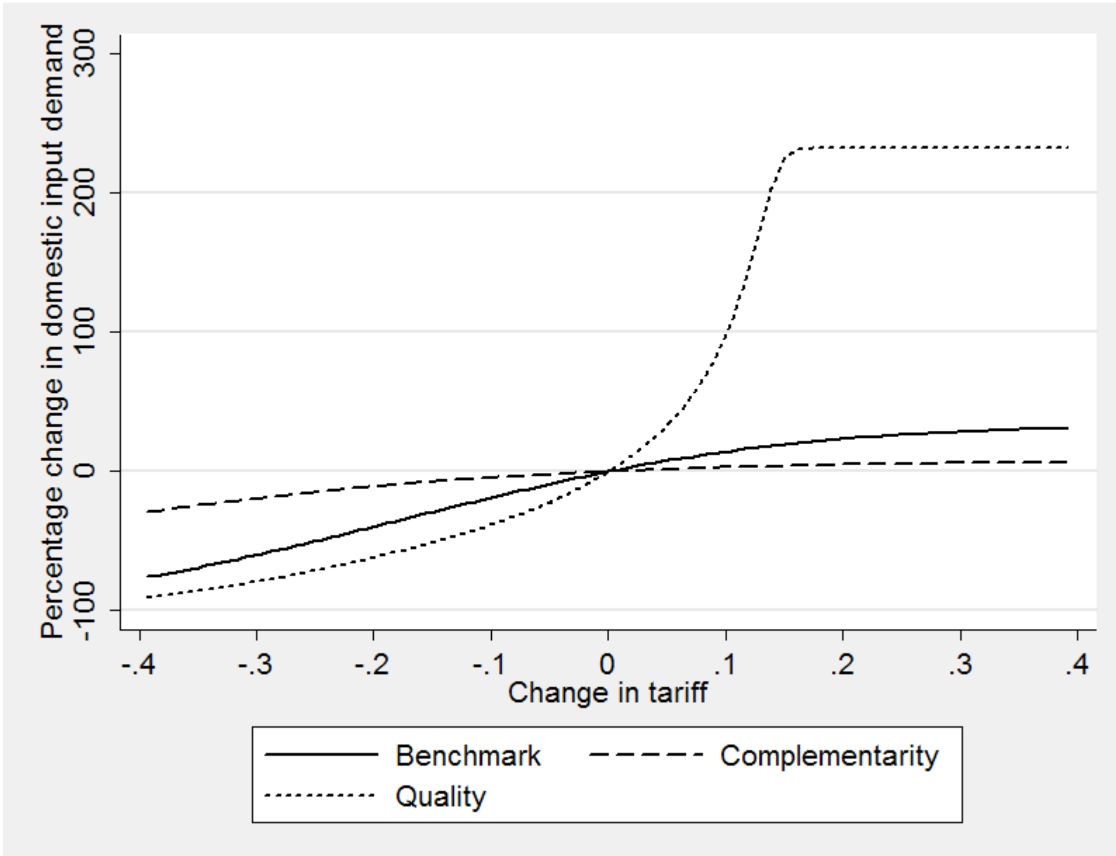

Figure 6: Tariff changes and domestic input demand

6.1 Heterogeneity in the benefit of imports

Our estimates assume that the benefits of importing, measured by B and θ, are identical across firms. Unfortunately, heterogeneity in these parameters may lead to endogeneity problems and possibly bias our estimates. Intuitively, firms who enjoy greater benefits from importing would both have higher productivity and import more. This omitted variable bias can introduce a spurious positive correlation between imports and productivity.

To assess the quantitative importance of this heterogeneity, consider the general version of our model where firms earn different quality gains from importing, so that B varies across firms, but the elasticity of substitution θ is constant. We show in the appendix that in this case the productivity of firm j can be written as

$$\text{TFP}_j = \frac{1}{\theta - 1} \ln\left(\frac{s_j}{1 - s_j}\right) G_j + \omega.$$

This equation shows that comparing across firms, productivity is an increasing, *convex* function of the import share s_j. Intuitively, when the quality gains from importing B vary across firms, the marginal gain from a unit increase in imports is higher for firms that already import a large amount, as these are the firms for whom imports are most beneficial. When B is constant across firms, the marginal gain from importing is also constant, and hence the relationship between productivity and the import share is linear.

This observation can be used to test for heterogeneity in B across firms. Figure 7 plots estimated productivity across firms against the share of imports in total cost. Estimated productivity is defined as $y - \hat{\alpha}k - \hat{\beta}l - \hat{\gamma}m$, where $\hat{\alpha}$, $\hat{\beta}$, $\hat{\gamma}$ are the estimated coefficients. The solid line is the result of a nonparametric regression used to capture any nonlinearities in the relationship between productivity and the import share.[11] The estimated relationship is remarkably linear: there is no evidence of a convex relationship between productivity and imports. This suggests that heterogeneity in B across firms is not a first-order source of variation in import patterns.

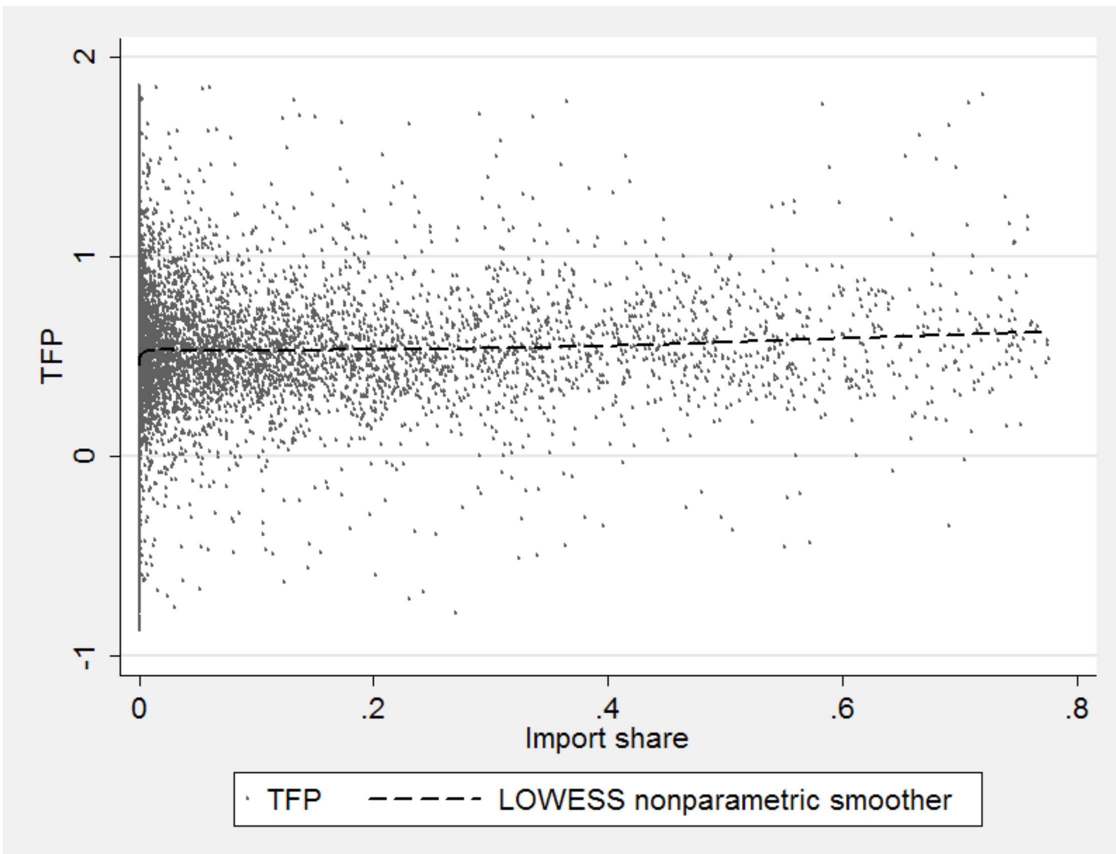

Figure 7: Productivity is linear in imports

[11]The nonparametric estimate is a locally weighted smooth regression (LOWESS) with a low bandwidth.

6.2 Heterogeneity across products

A second possible source of heterogeneity is that the parameter B_i may vary across products. For example, the productivity advantage of foreign goods may be much higher for computers than for copper wire. Similarly, products may have different elasticities of substitution, θ: e.g., computers may be more difficult to substitute than copper wire. This heterogeneity in parameters implies that different products have different a_i and s_i values. In this case, our procedure still consistently estimated the averages $\bar{a} = \sum_i \gamma_i a_i/\gamma$ and $\bar{s} = \sum_i \gamma_i a_i/\gamma$. The structural parameters B and θ are, however, a nonlinear function of a and s, and hence their estimates will no longer be consistent.

To get a sense of the magnitude of this bias, suppose that a fraction λ of products represent no productivity gains and have $a_i = 0$. This would be the case if $B_i = 0$, e.g., if the intermediate input is non-traded. Firms will not buy any of these products, hence they have $s_i = 0$. The rest of the products are assumed to bring equal productivity gains a. The estimated average productivity gains and average market shares are now simply $\bar{a} = (1-\lambda)a$ and $\bar{s} = (1-\lambda)s$. For any assumed value of λ, we can thus recover the underlying a and s parameters for traded products; and using equations (17) and (18), we can back out the relative quality B and the elasticity of substitution θ.

Table 5 presents the estimates of the structural parameters for different values of λ. Note that our data restricts λ to be at most 0.41, because we do observe firms who import 59% of their intermediate goods. The first row is the benchmark case, where all products are assumed to be traded. As the share of nontraded products increases, so does the inferred import share of traded products. For example, if $\lambda = 0.1$ and we see that the best firms import 59 percent of *all* their products, these imports represent 65 percent of *tradable* products. In our model, a high import share implies high elasticity of substitution; thus the estimate of θ is also increasing in the share of nontraded goods, and hence the importance of the imperfect substitution channel is reduced.[12] However, the results in Table 5 show that even with substantial heterogeneity, the imperfect substitution channel is responsible for a substantial share of the productivity gain, and hence our qualitative conclusions are unaffected.

6.3 Homogeneous and differentiated goods

When products differ not only in B_i but also in θ_i, we can no longer identify our structural parameters from the production function and the import demand equation alone. To see why, recall that χ_i is an indicator variable for whether the firm imports product i, and write

[12] More generally, it can be shown using equations (17) and (18) that heterogeneity in a always results in underestimating the elasticity θ.

Table 5: Estimates of B and θ with nontraded products

Share of nontraded inputs	**Import share of traded inputs**	**Relative quality**	**Elasticity of substitution**	**Percentage of gain from complementarity**
0.0	0.586	1.055	7.5	61%
0.1	0.651	1.094	8.0	41%
0.2	0.733	1.139	8.8	24%

Notes: This table reports estimates of the structural parameters under different assumptions about the share of nontraded inputs. See text for details.

the two equations as

$$y = \alpha k + \beta l + \gamma m + \sum_{i=1}^{N} \delta_i \chi_i + \omega,$$

$$X_F/Y = \sum_{i=1}^{N} \mu_i \chi_i + u,$$

where $\delta_i = \gamma_i a_i$ and $\mu_i = \gamma_i s_i$. Clearly, the $2N$ estimated coefficients of δ_i and μ_i are insufficient to identify all $3N$ parameters (a_i, s_i, γ_i). [13,14]

One approach is to proceed by assuming the additional restriction that a_i and s_i are small. In this case, we can use the first-order approximation of a_i around $s_i = 0$,

$$a_i \approx s_i/(\theta_i - 1),$$

which captures the idea that importing a given amount of input has stronger productivity gains when the elasticity of substitution is low. Now we can write $\delta_i \approx \mu_i/(\theta_i - 1)$, and then θ_i can be estimated as the ratio of δ_i and μ_i, even if the a_i, s_i and γ_i are themselves unidentified.

To illustrate this approach, we implement it using the classification of products into homogenous and differentiated categories due to Rauch (1999). We estimate δ and μ separately for these two classes. Table 6 presents the results. As column 2 shows, differentiated products have a bigger impact on productivity, $\delta_D > \delta_H$. This may be either because they are

[13] This was not an issue when all products had the same θ ($N - 1$ restrictions), because we could make use of the additional restriction that $\sum_{i=1}^{n} \gamma_i = \gamma$.

[14] One solution could be to use independent information on γ_is, the Cobb–Douglas shares of intermediates, for example, from input–output accounts. However, input–output accounts have broader level of aggregation than our 4-digit import statistics.

more important in production (higher γ) or because of the importance of complementarities – this coefficient does not yet identify which. Recovering the elasticity of substitution as $\hat{\theta}_i = 1 + \hat{\mu}_i/\hat{\delta}_i$, we find that differentiated goods have an elasticity of substitution of 5.6, whereas homogeneous goods have a significantly larger elasticity of 9.4. These results are consistent with the Rauch classification, as well with the estimates of Broda and Weinstein (2006).

Table 6: Homogeneous and differentiated products

	Homogeneous	**Differentiated**
Capital	0.021 (0.075)	
Labor	0.206 (0.010)	
Materials	0.782 (0.011)	
Productivity gain (a*gamma)	0.070 (0.023)	0.190 (0.073)
Import demand (s*gamma)	0.590 (0.020)	0.885 (0.193)
Elasticity (theta)	9.43 (2.35)	5.66 (0.38)
Year/industry fixed effects	Yes	
Full controls	Yes	
Number of observations	125,562	
Number of firms	32,067	
Tests:	p-value	
theta(H)>theta(D) (1-sided)	0.03	

Notes: Table reports the import parameters estimated separately for homogeneous and differentiated products. Product are split according to the classification by Rauch (1999). Other production function parameters are kept fixed. Bootstrapped standard errors (in parantheses) are clustered by firm. We proxy productivity by investment, controlling for export share, industry, year, and ownership category. Instruments include capital, lagged capital, and lagged imported capital. See text for details.

7 Conclusion

This paper explored the effect of imports on productivity in a panel of Hungarian firms. We structurally estimated a model of heterogenous importers, and found that imported inputs have large effects, much of which can be attributed to their complementarities with other goods in the production process. We also conducted counterfactual policy experiments, which demonstrated the importance of the extensive margin for new import varieties, and showed that tariff cuts have nonlinear effects on productivity.

We conclude by highlighting two areas where further work would be useful. First, combining a structural approach with data on plausibly exogenous policy changes would share the benefits of a formal model and straightforward identification. Second, our analysis ignores the imports of capital goods. Given recent work by Caselli and Wilson (2004), a complete understanding of imports should take into account the embedded technology in equipment as well.

References

Amiti, M. and Konings, J. (2007). Trade liberalization, intermediate inputs and productivity: Evidence from Indonesia, *American Economic Review* **97**.

Barro, R. J. (1997). *Determinants of Economic Growth*, MIT Press, Cambridge MA.

Bernard, A. B., Jensen, J. B. and Schott, P. K. (2009). Importers, exporters, and multinationals: A portrait of firms in the U.S. that trade goods, in T. Dunne, J.B. Jensen and M.J. Roberts (eds.), Producer Dynamics: New Evidence from Micro Data (University of Chicago Press).

Broda, C. and Weinstein, D. (2006). Globalization and the gains from variety, *Quarterly Journal of Economics* **121**(2).

Broda, C., Greenfield, J. and Weinstein, D. (2006). From groundnuts to globalization: A structural estimate of trade and growth, Working paper. Columbia University.

Caselli, F. and Wilson, D. (2004). Importing technology, *Journal of Monetary Economics* .

Coe, D. T. and Helpman, E. (1995). International R&D spillovers, *European Economic Review* **39**(5): 859–887.

Eaton, J., Kortum, S. S. and Kramarz, F. (2004). Dissecting trade: Firms, industries, and export destinations, *NBER Working Paper* **10344**.

Ethier, W. J. (1982). National and international returns to scale in the modern theory of international trade, *American Economic Review* **72**: 389–405.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices, *American Economic Review* **84**.

Frankel, J. A. and Romer, D. (1999). Does trade cause growth?, *American Economic Review* **89**: 379–399.

Goldberg, P., Khandelwal, A., Pavcnik, N. and Topalova, P. (2008). Imported intermediate inputs and domestic product growth: Evidence from india. Working paper, Princeton University, Columbia Business School, Dartmouth College and IMF.

Grossman, G. M. and Helpman, E. (1991). *Innovation and Growth in the Global Economy*, MIT Press, Cambridge MA.

Hallak, J. C. and Levinsohn, J. A. (2008). Fooling ourselves: Evaluating the globalization and growth debate. In: The Future of Globalization: Explorations in Light of Recent Turbulence.

Harrigan, J. and Choi, E. K. (eds) (2003). *Handbook of International Economics*, Blackwell.

Hirschman, A. O. (1958). *The Strategy of Economic Development*, Yale University Press.

Hummels, D. and Lugovskyy, V. (2004). Trade in ideal varieties: Theory and evidence, Working paper. Purdue University.

Jones, C. I. (2009). Intermediate goods and weak links: A theory of economic development. Working paper, Stanford GSB.

Kasahara, H. and Rodrigue, J. (2008). Does the use of imported intermediates increase productivity? plant-level evidence.

Keller, W. (2004). International technology diffusion, *Journal of Economic Literature* .

Konings, J. and Vandenbussche, H. (2008). Heterogenous responses of firms to trade protection, *Journal of International Economics* **76**.

Kremer, M. (1993). The o-ring theory of economic development, *Quarterly Journal of Economics* .

Krugman, P. (1979). Increasing returns, monopolistic competition, and international trade, *Journal of International Economics* **9**: 469–479.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity, *Econometrica* **71**: 1695–1725.

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

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry, *Econometrica* **64**: 1263–1297.

Rauch, J. E. (1999). Networks versus markets in international trade, *Journal of International Economics* **48**(1): 7–35.

Romer, P. (1987). Growth based on increasing returns due to specialization, *American Economic Review* **77**(2): 56–62.

Tybout, J. R. (2003). *Plant- and Firm-Level Evidence on "New" Trade Theories*, in Harrigan and Choi (2003).

Appendix

A Estimating the production function

The starting point is the assumption that, conditional on the observable state variable k, investment I is a monotonic function of productivity,

$$I_{jt} = f(\omega_{jt}, k_{jt}), \tag{19}$$

where j denotes firms and t stands for time. Assuming that (19) is invertible in its first argument yields

$$\omega_{jt} = g(I_{jt}, k_{jt}).$$

Combining this with (14) leads to

$$y_{jt} = \alpha k_{jt} + \beta l_{jt} + \gamma m_{jt} + \delta G_{njt} + g(I_{jt}, k_{jt}) + \varepsilon_{jt}, \tag{20}$$

where ε_{jt} is the part of productivity that is not observable to the firm and hence orthogonal to firm decisions. Because of this, estimating (20) in ordinary least squares yields consistent estimates of all coefficients, except for α which is not identified. We can now use the coefficients γ and δ to obtain an estimate of $\bar{a} = \delta/\gamma$.

Following Olley and Pakes, we deal with can deal with the unobserved function $g(I_{jt}, k_{jt})$, by including nonparametric functions of I_{jt} and k_{jt}.[15] The next step is to obtain a consistent estimate of α, which is not identified in (20). As Olley and Pakes, we assume that unobserved productivity ω follows a first-order Markov process, which for simplicity we approximate with an AR(1) that has autocorrelation ρ:

$$\mathrm{E}_t\, \omega_{j,t+1} = \rho \omega_{jt}.$$

For any given α and ρ, we can subtract ρ times $\omega_{j,t-1}$ from the current estimated ω_{jt} to obtain productivity innovations

$$u_{jt} \equiv [y_{jt} - \alpha k_{jt} - \beta l_{jt} - \gamma m_{jt} - \delta G_{jt}] - \rho g(I_{j,t-1}, k_{j,t-1}).$$

These innovations are orthogonal to all information available at time $t-1$, i.e., $\mathrm{E}(u_{jt}|\mathbf{Z}_{j,t-1}) = 0$. This orthogonality condition can be used in a generalized method of moments estimation with instrument set Z. We use as instruments current and lagged capital, lagged employment, lagged material cost, lagged number of products, and the lagged import share. An additional problem is that we do not observe u_{jt} for exiting firms, so we can only calculate $\mathrm{E}(u_{jt}|\mathbf{Z}_{j,t-1})$ conditional on firm survival. To correct for this bias, we first obtain propensity scores for exit by running a linear probability regression on current and lagged capital

[15] In the implementation we use third-order polynomials of I, k, and other control variables.

and other controls, and then control for this propensity score nonparametrically. Finally, we estimate $\rho$from a grid search over $[0, 1]$ seeking to minimize the weighted squared sum of moments. The J-test of overidentification is not rejected in any of our specifications, verifying that productivity innovations are indeed orthogonal to all of the instruments. This procedure yields estimates of the parameters α, β, γ, δ, and hence $\bar{a} = \delta/\gamma$ as well.

B Heterogenous effects across firms

In this appendix we consider the general version of our model where firms earn different quality gains from importing, so that B varies across firms. Then the productivity of firm j can be written as

$$\text{TFP}_j = a_j G_j + \omega = \frac{1}{\theta - 1} \ln\left(1 + B_j^{\theta-1}\right) G_j + \omega,$$

which is increasing in B_j. The optimal import share of this firm is

$$s_j = \frac{B_j^{\theta-1}}{1 + B_j^{\theta-1}},$$

which is also increasing in B_j. Substituting out B_j from the import demand equation, we can express firm productivity as a function of the import share as

$$\text{TFP}_j = \frac{1}{\theta - 1} \ln\left(\frac{s_j}{1 - s_j}\right) G_j + \omega.$$

This equation shows that productivity is an increasing, *convex* function of the import share s_j. Intuitively, the marginal gain from a unit increase in imports is higher for firms that already import a large amount, as these are the firms for whom imports yield the highest quality gains. The convex relationship between import share and productivity arises only because of heterogeneity in B: when B is constant across firms, the marginal gain from importing is also constant, and hence the relationship between productivity and the import share is linear.