Imported Inputs and Productivity*

László Halpern

Institute of Economics, Hungarian Academy of Sciences and CEPR

Miklós Koren

Central European University, IEHAS and CEPR

Adam Szeidl

Central European University, UC-Berkeley, NBER and CEPR

September 2011

Abstract

We estimate a model of importers in Hungarian micro data and conduct counterfactual policy analysis to investigate the effect of imports on productivity. We find that importing all foreign varieties would increase firm productivity by 12 percent, almost two-fifths of which is due to imperfect substitution between foreign and domestic goods. The effectiveness of import use is higher for foreign firms and increases when a firm becomes foreign-owned. Our estimates imply that during 1993-2002 one-third of the productivity growth in Hungary was due to imported inputs. Simulations show that the productivity gain from a tariff cut is largest when the economy has many importers and many foreign firms, implying policy complementarities between tariff cuts, dismantling non-tariff barriers, and FDI liberalization.

1 Introduction

Understanding the link between international trade and aggregate productivity is one of the major challenges in international economics. To learn more about this link at the microeconomic level, a recent literature explores the effect of imported inputs—which constitute the majority of world trade—on firm productivity. Studies show that improved access to foreign inputs has increased firm productivity in several countries, including Indonesia (Amiti and Konings 2007), Chile (Kasahara and Rodrigue 2008) and India (Topalova and Khandelwal 2011).[1] The next step in this research agenda is to investigate the underlying mechanism through which imports increase productivity. As Hallak and Levinsohn (2008) emphasize, understanding which firms gain most, through what channel, and how the effect depends on the economic environment, are important for evaluating the welfare and redistributive implications of trade policies.

To explore these questions, in this paper we estimate a structural model of importer firms in Hungarian firm level data, and conduct counterfactual policy analysis in our estimated economy. Our starting point is a unique dataset that contains detailed information on imported goods for essentially all Hungarian manufacturing firms during 1992-2003. Motivated by stylized facts in these data, we formulate a model of firms who use differentiated inputs to produce a final good and must pay a fixed cost for each variety they choose to import. Imported inputs affect firm productivity through two distinct channels emphasized in the literature: they have a potentially higher price-adjusted quality as in quality-ladder models, and they imperfectly substitute domestic inputs as in product-variety models.[2] Because of these forces, firm productivity increases in the number of varieties imported. Our model also permits rich heterogeneity across products and firms.

In the first half of the paper we estimate this model in micro data. In doing so, we face the key empirical challenge that imports are chosen endogenously by the firm. We deal with this identification problem using a structural approach which exploits the product-level nature of the data. Our model implies a firm-level production function in which output depends on the usual factors of production as well as a new term related to the number of imported varieties. To estimate this production function, we follow Olley and Pakes (1996) in nonparametrically controlling for firm investment and other state variables, which pick up the unobserved component of productivity under standard assumptions. The import effect

---

[1]Results are conflicting for Brazil: Schor (2004) estimates a positive effect while Muendler (2004) finds no effect of imported inputs on productivity.

[2]For quality-ladder models see Aghion and Howitt (1992) or Grossman and Helpman (1991). Variety effects are introduced in Ethier (1982).

is then identified from residual variation in the number of imported varieties. Intuitively, we estimate the difference in output between two firms that have the same productivity, but differ in the number of varieties they choose to import—which, according to our model, happens because they face a different fixed cost of importing. We find substantial gains from imported inputs: in the baseline specification, increasing the fraction of tradeable goods imported by a firm from 0 to 100 percent would increase productivity by 12 percent.

We then turn to decompose the import effect into differences in quality and imperfect substitution. We first note that for a given productivity gain from importing a particular good, the degree of substitution governs a firm's expenditure share of foreign versus domestic purchases. For example, when foreign and domestic inputs are close to perfect substitutes, even if the productivity gain from imports is small the import share should be high.[3] Based on this idea, we then infer the relative magnitude of the two channels by comparing the expenditure share of imports for firms who differ in the number of imported varieties. While quality effects are important, we find that combining imperfectly substitutable foreign and domestic varieties is responsible for a substantial part—almost 40 percent—of the productivity gain from imports. This finding parallels the evidence in Goldberg, Khandelwal, Pavcnik and Topalova (2009) that combining foreign and domestic varieties increased firms' product scope in India; and also the theoretical arguments of Hirschman (1958), Kremer (1993) and Jones (2011) that complementarities, which amplify differences in input quality, may help explain large cross-country income differences.

We next explore whether the benefits from importing differ between foreign and domestic firms. Intuitively, because they have know-how about foreign markets and can access cheap suppliers abroad, foreign companies may gain more from spending on imports. This is an important possibility because foreign firms have played a very significant role in Hungary: during 1992-2003, their sales share in our data increased from 21 percent to 77 percent. When we re-estimate our model allowing for differences in the efficiency of import use by ownership status, we find that foreign companies benefit about 27 percent more than domestic ones from each dollar they spend on imports. To explore whether this large difference is directly caused by foreign ownership and not by selection, we use an event study in the subsample of initially domestic companies which were subsequently purchased by foreign investors. We find that the effectiveness of import use increases by 13.5 percent after a company changes owners (p-value of 9.9%), providing suggestive evidence—in a relatively small sample—that

---

[3]This link between import demand and the role of complementarities is also exploited by Feenstra (1994), Broda and Weinstein (2006) and Broda, Greenfield and Weinstein (2006) in country level data.

at least part of the advantage of foreign firms is *caused* by them being foreign-owned, and implying a complementarity between foreign presence and importing.

In the second half of the paper, we study the economic and policy implications of our estimates in two applications. We first quantify the contribution of imports to productivity growth in Hungary during 1993-2002. Our estimates imply a productivity gain of 15 percent in the Hungarian manufacturing sector, of which 6 percentage points, more than a third, can be attributed to import-related mechanisms. Approximately half of these import-related gains are due to the increased volume and number of imported inputs, while the other half is the result of increased foreign ownership in combination with foreign firms being better at using imports. These findings show both that imports contributed substantially to economic growth in Hungary, and the large aggregate effect of the complementarity between foreign presence and importing.

In our second application we use simulations in the estimated economy to explore the productivity implications of tariff policies. Intuitively, a tariff cut, by reducing the cost of and increasing access to foreign inputs, should raise both firm-level and aggregate productivity. Our main result is that the size of the aggregate productivity gain depends on two key features of the environment: (1) the initial import participation of producers; (2) the degree of foreign presence. Perhaps surprisingly, a higher initial import participation—more firms importing more kinds of products—implies larger gains from a tariff cut. This is because the set of inputs whose costs are affected is larger, and hence firms save more with the tariff cut. This logic also implies that the productivity effect of tariff cuts is convex: larger cuts have a more-than-proportional effect because they also increase the set of imported goods. In turn, foreign presence matters because, as we have shown, foreign-owned firms are better in using imports.

These patterns lead to complementarities between different liberalization policies. For example, our simulations show that tariff cuts increase productivity more when the fixed costs of importing—such as licensing or other non-tariff-barriers—are also reduced. Intuitively, a lower fixed cost makes a larger set of goods available, expending the scope for cost-savings from the tariff cut. Because foreign firms are more effective in using imports, a similar complementarity exists between tariff cuts and FDI liberalization. While a careful analysis of other economies is beyond the scope of this paper, these complementarities seem broadly consistent with the liberalization experience in the early 1990s in India. Consistent with the first complementarity, tariff cuts in India, which were accompanied by dismantling substantial non-tariff barriers, lead to rapid growth in new imported varieties (Goldberg et al. 2009) and a large increase in firm productivity (Topalova and Khandelwal 2011). And

3

consistent with the second complementarity, these effects were stronger in industries with higher FDI liberalization (Topalova and Khandelwal 2011).

Our tariff experiment also highlights the differential policy implications of the quality and imperfect substitution mechanisms. We show that the demand for domestic inputs is more sensitive to tariffs when the benefit of imports comes from quality differences than when it comes from imperfect substitution. This is intuitive: when foreign goods are close to perfect substitutes, even a small price change can bring about large import substitution. Because our estimates assign a significant role to imperfect substitution, we obtain a relatively inelastic demand curve for domestic inputs. In addition, losses to domestic input suppliers caused by the tariff cut are partially offset by increased demand for their products due to higher firm productivity.[4] One implication of these two forces is that redistributive losses due to import substitution may not be very large in practice. A broader lesson from our counterfactual analysis is that identifying the mechanism through which imports affect productivity helps evaluate different trade policies.

Besides the papers cited above, we build on a growing empirical literature exploring firm behavior in international markets, reviewed in Bernard, Jensen, Redding and Schott (2007). Tybout (2003) summarizes earlier plant and firm level empirical work testing theories of international trade. Our structural approach parallels Das, Roberts and Tybout (2007) who study export subsidies, and Kasahara and Lapham (2008) who investigate the link between exports and imports, by estimating optimizing models. Our basic theoretical framework also builds on work by Ethier (1979) and Markusen (1989) who develop models connecting imported inputs and productivity.

The rest of this paper is organized as follows. Section 2 describes our data and documents stylized facts about importers in Hungary. Building on these facts, in Section 3 we develop a simple model of importer-producers. Section 4 describes the estimation procedure and results. In Section 5 we use these estimates to conduct counterfactual analysis. We discuss some caveats with our approach in the concluding Section 6.

---

[4]This logic parallels Grossman and Rossi-Hansberg (2008) who argue that offshoring can sometimes— surprisingly—increase domestic labor demand due to the increase in output.

2 Data

2.1 Data and sample definition

Our panel of essentially all Hungarian manufacturing firms during 1992-2003 is created by merging trade and balance sheets data. Annual exports and imports for all firms, disaggregated by products at the 6-digit Harmonized System (HS) level (5,200 product categories), come from the Hungarian Customs Statistics. Because the 6-digit classification is noisy, we aggregate the data to the 4-digit level (1,300 categories), and we use the terms "product" and "good" to refer to a HS4 category in the rest of the paper.[5] Firms' balance sheets and profit and loss statements come from the Hungarian Tax Authority for 1992-1999, and from the Hungarian Statistics Office for 2000-2003. The data for 1992-1999 contain all firms which are required to file a balance sheet with the tax authority, i.e., all but the smallest companies—the main category of omitted firms are individual entrepreneurs without employees. The data for 2000-2003 includes all firms with at least 20 employees and a random sample of firms with 5-20 employees. We thus lose some firms in 2000, which, however, constitute a relatively small share of output: during 1992-1999, firms with no more than 20 workers were responsible for less than 7.5 percent of total sales. We classify a firm to be in the manufacturing sector if it reports manufacturing as a primary activity for at least half of its lifetime in the data, and exclude all other firms. We merge the trade data and the balance sheet data using firms' tax identifiers.

While our data contain product level information on *imported* input use, a limitation is that we do not have corresponding product-level data for *domestic* input purchases, because balance sheets only measure total spending on intermediate goods. We will rely on our structural model and on input-output tables to work around this data issue. A second limitation is that we do not observe at the firm level import purchases from domestic wholesalers such as export-import companies. We can, however, measure the role of such indirect imports for the economy as a whole. We find that the total value of intermediate imports by wholesalers and retailers—as computed from our data—is about 2 percent of total intermediate input use by all firms in all sectors of the Hungarian economy. This result suggests that the role of intermediation is relatively small for intermediate goods, and due to lack of additional data we ignore it below.

Processing trade. An important source of measurement error in our data is that some firms engage in processing trade. In exchange for a fee, these firms import, process and

---

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

re-export intermediate goods which remain the property of a foreign party throughout. Because the processing firm does not own, purchase or sell the underlying goods, processing trade is not recorded on the firm's balance sheet, even though it appears in our trade data. This inconsistency creates measurement problems: in several observations, the value of imported intermediate inputs, as measured by customs, *exceeds* total spending on intermediate goods from the balance sheet. Similarly, exports measured in the customs data are often substantially higher than exports from the balance sheet.

To deal with this reporting discrepancy, we measure processing trade performed by a firm as the difference between customs exports and balance sheet exports (when positive). We classify a firm as "processer" in a given year if the ratio of processing trade to balance sheet sales exceeds 2.5 percent, which is approximately the median across observations where this ratio is positive. About 10 percent of our observations are classified as processers. To obtain measures which reflect the underlying economic activity rather than accounting rules, we then adjust, for all firms, sales and total intermediate spending from the balance sheet by adding our measure of processing trade.

Following these changes we create two data samples for our analysis. Our *main sample* is defined by excluding all firm-year observations in which the firm is classified as a processer. We also define a *firm-level sample* which is obtained by fully excluding firms which are processers for more than half of the years they are in our sample. The reason for the exclusions is that our adjustment for processing likely introduces considerable noise.[6] Because it has more observations, unless otherwise noted we will use the main sample in our analysis. The benefit of the firm-level sample is that, because it does not permit changes in the set of firms over time due to changes in processing activity, it better reflects aggregate trends in the data. After all exclusions, 130,027 observations remain in our main sample.

Input-output table. Besides the firm level data, we also make use of a sectoral input-output table at the 2-digit International Standard Industrial Classification (ISIC, revision 3) for the year 2000, which comes from the Hungarian Statistics Office.

2.2 Summary statistics and stylized facts

We document three basic facts about firms' import behavior in the data, which will guide the specification of our formal model in Section 3.

---

[6]While we believe the exclusions are justified on prior grounds, keeping these firms in the sample and including an indicator for processers in all empirical specifications does not affect our qualitative results.

Table 1: Descriptive statistics

	Full sample	Non-importers	Importers
Employment	48.390	17.500	98.040
Sales (thousand USD)	2,908	392	6,952
Capital per worker (thousand USD)	12.990	10.700	16.660
Sales per worker (thousand USD)	49.410	35.340	72.030
Material share in output	0.662	0.636	0.705
Exporter indicator	0.351	0.149	0.675
Export share in output	0.108	0.047	0.205
Importer indicator	0.383		
Import share in materials	0.105		0.274
Number of imported products (HS4)	4.660		12.160
Foreign owned	0.157	0.071	0.296
State owned	0.027	0.020	0.037
Number of observations	130,027	80,162	49,865
Number of firms	27,074	21,361	13,556

Notes: Tables entries are means unless otherwise noted. Column 1 is based on the full sample defined in Section 2.1. Column 2 is computed for firm-years in which the firm does not import, and column 3 is computed for firm-years in which the firm does import. The number of firms in columns 2 and 3 add up to more than the total number of firms because of firms that switch importer status. USD values are in 1991 dollars.

Fact 1. There is substantial heterogeneity in the import patterns of firms. 61 percent of firms do not import at all; larger and foreign-owned companies are more likely to be importers.

This fact can be seen by comparing across columns in Table 1, which presents summary statistics for several key variables in our main sample separately for importing and non-importing firms. Importers employ about 6 times as many workers and sell about 18 times as much as non-importers. Using the classification that a firm is foreign-owned in a given year if foreigners have majority ownership, importers are also more frequently foreign owned and more likely to export.[7] There is also substantial heterogeneity within importers in the number of products they import. Regressing the log number of imported products on log employment and foreign ownership shows that doubling firm size is associated with a 26 percent increase in the number of imported products, and, conditional on size, foreign firms import 144 percent more products than domestic firms. The patterns shown here are

---

[7]Firm level evidence from other countries shows similar patterns: for example, Bernard, Jensen and Schott (2009) document that "globally engaged firms" in the U.S. are superior along a number of dimensions.

consistent with a model in which entry in import markets entails a fixed cost. Larger or more productive firms profit more from a given product and hence find it easier to overcome the fixed cost. Similarly, foreign firms may have lower fixed or variable costs of importing and hence purchase more foreign varieties.

Fact 2. Import spending is concentrated on a few core products; firms spend little on their remaining imports.

To document this fact, for each firm, we order imported products by the share of total import spending allocated to them. Using this ranking, among firms importing five or more products, the average spending share (out of total import spending) of the highest-ranked product is 54 percent. Thus, on average, firms spend more than half of their import budget on a single product. In contrast, the average spending share of the fifth-highest ranked product is only 3.5 percent. This substantial heterogeneity across goods may be important for evaluating the productivity gain from importing new products.

Fact 3. The extensive margin plays a large role in explaining both the aggregate trend and the firm level fluctuations in import growth.

Table 2, constructed from our firm-level sample, shows aggregate trends in firm imports over time. It is instructive to look at average annual growth as well as growth over the entire period of 1992-2003, reported in the last two rows. As column 3 shows, intermediate imports in the manufacturing sector grew by 23.4 percent per year on average, resulting in a total import growth of 657 percent between 1992 and 2003. On a year-to-year basis, almost 85 percent of import growth, 19.8 percent, is explained by the within-firm intensive margin, i.e., increased imports of goods that the firm was already importing in the previous year. Over the entire decade, however, the role of the extensive margin is much larger: imports of firms that did not import in 1992 and of products that were not imported in 1992 explain 84 percent of the growth in imports 1992-2003. The extensive margin also has a large yearly *gross* volume: on average, new importers and newly imported products contribute 10.8 percent, while exit and the firm and the product level contribute -7.1 percent to import growth per year.[8] These patterns call for an explicit model of the extensive margin of imports, and—given the high frequency of exit at both the firm and product level—suggest that imports likely entail per-period fixed costs.

---

[8]Due to the change in sample definition in 2000, we lose some importing firms in that year (see Section 2.1). These firms are also classified as stopping importers in the table. Because we only lose firms with 20 or fewer employees, the vast majority of which do not import, their effect on the volume-weighted numbers in the table is likely to be small.

Table 2: Import dynamics

Year	Total imports	Import growth	Contribution to import growth				
			Intensive margin	New importers	New products	Stopping importers	Dropped products
1992	1,717						
1993	2,036	18.6%	15.0%	15.5%	8.6%	-14.7%	-5.8%
1994	2,758	35.4%	28.8%	8.7%	7.4%	-4.6%	-4.8%
1995	3,503	27.0%	22.6%	5.9%	5.4%	-2.9%	-4.0%
1996	4,676	33.5%	12.4%	6.8%	23.4%	-4.3%	-4.8%
1997	6,291	34.6%	32.6%	2.9%	3.9%	-2.0%	-2.8%
1998	8,728	38.7%	37.3%	1.5%	3.0%	-0.8%	-2.2%
1999	11,272	29.1%	26.7%	2.7%	2.2%	-1.1%	-1.3%
2000	14,622	29.7%	28.2%	11.0%	1.5%	-9.8%	-1.1%
2001	18,199	24.5%	23.2%	1.4%	2.7%	-2.1%	-0.7%
2002	17,218	-5.4%	-2.2%	0.7%	1.0%	-4.2%	-0.8%
2003	15,759	-8.5%	-7.3%	1.3%	0.9%	-1.1%	-2.4%
Average		23.4%	19.8%	5.3%	5.5%	-4.3%	-2.8%
1992 to 2003		657.4%	152.9%	501.0%	54.3%	-38.8%	-12.0%

Notes: Total imports are in millions of 1991 USD. Columns 4-8 measure the percent increase in imports attributable to different mechanisms and sum to column 3 in each year. The intensive margin measures the contribution of (net) growth in imports of products that the firm also imported the previous period (which in all rows but the last is the previous year, and in the last row is 1992). New importers are firms that did not import in the previous period. New products are newly imported products of existing importers. Stopping importers and dropped products are defined analogously.

3 A Model of Firm Imports and Productivity

Motivated by the above stylized facts, we build a model of a firm which uses both domestic and imported intermediate goods for production. Section 4 will introduce the types of heterogeneity across firms that we allow in estimating this model.

3.1 Setup

Production technology. The output of the firm is given by the production function

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

where K and L denote capital and labor used in production, X_i denotes intermediate composite good i, and Ω is Hicks neutral total factor productivity (TFP). The Cobb-Douglas

weight γ_i measures the importance of good i for production. Motivated by Fact 2, we allow γ_i to be different for different goods i.[9] The total weight of all intermediate goods is $\gamma = \sum_i \gamma_i$.

Each intermediate good X_i is assembled from a combination of a foreign and a domestic variety:

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

where X_{iF} and X_{iH} are the quantity of foreign and domestic inputs, and θ is the elasticity of substitution. The prices of the domestic and foreign varieties are denoted p_{iH} and p_{iF}, and by adjusting units, we set $p_{iH} = 1$. The price-adjusted quality of the foreign input is then $A_i = B_i/p_{iF}$. Intuitively, A_i measures the advantage of a dollar spent on a foreign relative to a domestic variety.

To aid the estimation, we make several assumptions about intermediate inputs. In practice some goods used for production are non-tradeable, and hence are not imported. To allow for non-tradeables in a simple way, we assume that they coincide with the set of services, and assign an infinitely high foreign price and hence $A_i = 0$ to these inputs. Doing this will allow us to estimate the input share of non-tradeables from an input-output table. We also assume that the price-adjusted quality A_i of all tradeable goods used by the firm is the same $A_i = A > 0$. This assumption greatly simplifies our analysis and still allows us to estimate the *average* quality advantage of imports.[10] Finally, we order indices so that inputs $1, 2..., N_g$ represent tradeable goods, while the remaining $N_g + 1, ..., N$ inputs represent non-tradeable services. We also order tradeable goods by their production weight, so that $\gamma_1 \geq \gamma_2 \geq ... \geq \gamma_{N_g}$.

Motivated by stylized fact 3, we assume that the firm must pay a fixed cost F to access foreign markets, and that importing each variety has an additional fixed cost f which is constant across products within a firm. To make the model consistent with the high frequency of exit from import markets we assume that these costs are due every period.[11]

The firm sells its final good in a monopolistically competitive market, facing a firm-specific demand curve. For estimating the basic parameters, we require that the demand curve is sufficiently downward-sloping so that the solution to the firm's profit-maximization problem exists.[12] To obtain estimates of the fixed costs, and for the counterfactual analysis

---

[9]Hummels and Lugovskyy (2004) develop a related model where the marginal utility of additional varieties declines.

[10]We do allow price-adjusted quality A to vary across firms by industry, ownership status and over time in the estimation.

[11]In the estimation we also explore the case when F is a sunk cost.

[12]Denoting the (local) elasticity of demand by η, this is the case when $\alpha + \beta + \gamma < \eta/(\eta - 1)$ holds everywhere.

of Section 5.2, we impose a parametric isoelastic demand curve.

Discussion. Our production specification incorporates both the quality and variety gains from importing emphasized in the literature. Following Grossman and Helpman (1991), we interpret quality as the advantage in services provided by a good *relative to its cost.* The natural measure of the quality gain is therefore price-adjusted quality A, which can also be interpreted as the efficiency with which firms use imports, relative to domestic products, per dollar of spending. Imperfect substitution, i.e., the idea that combining foreign and domestic goods create gains that are greater than the sum of the parts, is measured by the elasticity of substitution θ. For θ finite, we always have $a > \log A$: imperfect substitution amplifies the gain from a higher-quality foreign input. Our setup thus allows for flexibility in the degree of substitution as well as heterogeneity across inputs while maintaining the tractability of the Cobb-Douglas model. As we show below, this framework also gets around a data limitation by generating estimating equations for output and imports that involve only imported, but not domestic product level input purchases.

Because we are primarily concerned with the importing behavior of a firm at a point in time, the model introduced here is static. We make additional, standard, assumptions about firm dynamics in Section 4.1 to aid the estimation of this model in panel data.

3.2   Firm behavior

To characterize firm behavior, we first compute the cost-saving when the firm uses also imports for a particular composite good. Given the normalization that the domestic price of all inputs is $p_{iH} = 1$, the effective price of a composite good X_i is $P_i = 1$ if the firm only uses the domestic input X_{iH}. If the foreign input is also used, the effective price can be found by solving the cost-minimization problem associated with (2):

$$P_i = \left[p_{iH} + (p_{iF}/B_i)^{1-\theta} \right]^{1/(1-\theta)} = \left[1 + A^{\theta-1} \right]^{1/(1-\theta)} < 1 \tag{3}$$

using the notation that $A_i = B_i/p_{iF}$ and our assumption that $A_i = A$ for all tradeable imported inputs. The (log) percentage reduction in the cost of the composite good i when imports are used is therefore

$$a = \frac{\ln \left[1 + A^{\theta-1} \right]}{\theta - 1}. \tag{4}$$

Parameter a measures the *per-product import gain* and hence is of central interest to us. Note that a incorporates the cost-savings created by both the quality and the imperfect-substitution channels: in particular, it is increasing in the price-adjusted quality A and decreasing in the degree of substitution θ. Because of imperfect substitution, for finite θ the

11

firm uses both domestic and foreign inputs, so that the optimal expenditure share of the foreign good in the total spending for variety i, which can be computed as

$$s = A^{\theta-1}/(1 + A^{\theta-1}) \tag{5}$$

satisfies $0 < s < 1$.

We next characterize the choice of which varieties to import. In this decision, the firm trades off the savings in marginal cost from using imports against the fixed costs of importing. Since the fixed cost f and the per-product gain a are the same across products, a firm which imports n products will choose to import those with the highest γ weight, i.e., products $i = 1, ..., n$. Let $\pi(n, F, f)$ denote profits if the firm imports these n goods given cost realizations f and F.[13] The optimal import decision of the firm is then

$$n = \arg \max_{\tilde{n}} \pi(\tilde{n}, F, f). \tag{6}$$

It is easy to see that n is weakly decreasing in f and F: firms with lower fixed costs import a greater number of varieties.

Once the firm sets n, the relative importance of imported inputs for production can be measured as the Cobb-Douglas share of composite goods which use imports relative to all intermediate inputs, which we denote

$$G(n) = \frac{\sum_{i=1}^{n} \gamma_i}{\sum_{i=1}^{N} \gamma_i} = \frac{\sum_{i=1}^{n} \gamma_i}{\gamma}. \tag{7}$$

Since $\gamma_1 \geq \gamma_2, ... \geq \gamma_g \geq 0$, the $G(\cdot)$ function is increasing and concave. Because the denominator includes the weights of both goods and non-traded services, the maximum of $G(\cdot)$, denoted $\bar{G} = G(N_g)$, equals the share of tradable goods in all intermediate inputs.

We now use $G(n)$ to connect the number of imported inputs to import demand and firm output. These two equations will form the basis for our empirical analysis. Denoting total expenditure on intermediate inputs by $M = \sum_i P_i X_i$ and total expenditure on foreign inputs by $M^F = \sum_{i=1}^{N} p_{iF} X_{iF}$, the spending share on imports can be written as

$$\frac{M^F}{M} = s \frac{\sum_{i=1}^{n} \gamma_i}{\gamma} = sG(n) \tag{8}$$

where s, defined in (5), is the optimal expenditure share of imports within a composite good i. This equation links the import share to the number of imported varieties n. Intuitively,

---

[13]For notational simplicity we suppress the dependence of π on other firm level variables such as k or ω. The profit function can be calculated using the demand curve for the firm's product, and we compute it explicitly in Appendix A for isoelastic demand.

firms that import a greater number of products n have a larger share of foreign goods in total intermediate spending. This effect is captured by $G(n)$ because the Cobb-Douglas production function implies that the input share of each imported variety i equals $s\gamma_i/\gamma$.

For an optimizing firm, the following expression—proved in Appendix A—expresses total output with the number of imported varieties:

$$y = \alpha k + \beta l + \gamma m + a\gamma G(n) + \omega \tag{9}$$

where the lowercase variables y, k, l, m and ω denote logs. This equation has the form of a standard production function: the first three terms on the right-hand side measure the contribution of capital, labor and intermediate inputs to total output, while the final term is the Hicks-neutral productivity shifter ω. The novelty lies in the fourth term, which represents the contribution of imports. The intuition for this term is straightforward: a firm which chooses to import n varieties will have a percentage cost reduction of a on the associated composite inputs, the total weight of which is $\sum_{i=1}^{n} \gamma_i$. This cost reduction maps into a corresponding increase in output for a given total spending on intermediate inputs.

A natural interpretation of equation (9) is that a firm's total factor productivity is given by $\phi = a\gamma G(n) + \omega$, i.e., the sum of the productivity gains from importing and a "residual productivity" term. This interpretation is correct in the sense that variation in ϕ measures differences in output for the same amount of resources *employed in the production process*; but it ignores the fact that importing also entails fixed costs which require resources. Thus ϕ is an (approximately) correct measure of productivity holding fixed all resources only when the fixed costs are small relative to the overall productivity gain, which—because importing reduces *marginal* costs but requires the payment of *fixed* costs—is more likely to be the case for larger firms who import multiple different products.[14] In the empirical analysis, we will compute an adjusted measure of productivity which also reflects the fixed costs of importing, and show that—because the bulk of production and importing is performed by mid-sized and large importers—the two measures generate essentially identical aggregate predictions. Hence in practice little is lost by treating ϕ as a measure of productivity, which is what we do below.

---

[14]More precisely, for the last product the firm chooses to import, the fixed cost should be approximately the same as the savings induced by importing that product. For every other—inframarginal—product that the firm chooses to import, the fixed cost of importing is strictly lower than the cost-saving from lower marginal costs, and this difference is increasing in firm size because larger firms gain more from a given reduction in marginal cost.

4   Estimation and Results

We now turn to estimating the model in firm level data. Section 4.1 develops our empirical strategy and Section 4.2 presents the results.

4.1   Empirical strategy

Assumptions about firms and products. Our theoretical model is static and applies to a single firm. We now state assumptions about heterogeneity and dynamics which allow us to estimate this model in panel data.

Our empirical strategy assumes that firms can be partitioned into different groups—for example by industry or ownership status—such that certain model parameters are identical within a group, but can potentially differ between groups. The basic parameters α, β, γ_i, A and θ are always held constant within a group, and we also make—standard—homogeneity assumptions about dynamics to ensure that the Olley-Pakes procedure yields consistent estimates within each group. Specifically, we write $\omega_{jt} = \omega_{jt}^{obs} + \varepsilon_{jt}$ for a firm j in period t, where ω_{jt}^{obs} is observable to the firm at the beginning of period t and ε_{jt} is a mean-zero innovation, independent of other firm variables and identically distributed across firms in a group and over time. We assume that investment I_{jt}, which is set before ε_{jt} is realized, can be written as $I_{jt} = \xi(\omega_{jt}^{obs}, z_{jt})$ where z_{jt} is a vector of observable firm state variables which always includes capital k_{jt} and calendar time t, and the ξ function is the same across firms in a group and is strictly increasing in its first argument. In addition, ω^{obs} follows a first-order Markov process with the same dynamics across firms in a group. These assumptions are common in the productivity literature. We also assume that the fixed costs F_{jt} and f_{jt} are drawn after the investment decision is made, independently across firms and over time, from a distribution that may depend on z_{jt} and ω_{jt}^{obs}.

This framework allows for considerable heterogeneity within a group: firms can differ in their productivity, factor use, foreign and domestic intermediate input use, and also in their realized fixed costs. We also permit heterogeneity across products through the γ_i parameters. The requirement that the γ_i—essentially, the $G(\cdot)$ function— are the same within a group means that additional varieties decline in importance identically across companies, so that all firms in a group have the same production structure. Importantly, it *does not* mean that firms use the same goods in production, or that goods have the same production weight. For example, γ_1, the share of the most important input, is the same for all firms; but this share can be different from γ_2, and also, the identity of the most important good can vary across firms.

Estimating equations. We now use our assumptions to convert the production function (9) and the import demand equation (8) into estimable forms by substituting out $G(n)$ and the productivity term ω^{obs}. Because the shape of $G(\cdot)$ is unknown, we approximate it with an exponential function $G(n) = \bar{G}\left[1 - \exp(-\lambda n)\right]$ whose parameters $\bar{G} < 1$ and $\lambda > 0$ we estimate.[15] To control for the productivity shifter ω^{obs}, we follow Olley and Pakes (1996) and invert the monotonically increasing investment function ξ to get

$$\omega_{jt}^{obs} = h(I_{jt}, z_{jt})$$

with an unknown h "control" function, which is the same across all firms in a group.[16] Substituting the expressions for $G(\cdot)$ and ω^{obs} into (9) yields our first estimating equation

$$y_{jt} = \alpha \cdot k_{jt} + \beta \cdot l_{jt} + \gamma \cdot m_{jt} + \delta \cdot \left[1 - \exp(-\lambda n_{jt})\right] + h(I_{jt}, k_{jt}) + \varepsilon_{jt}, \qquad (10)$$

where $\delta = a\gamma\bar{G}$, and given our assumptions, ε_{jt} is independent of all variables on the right hand side.

Our second estimating equation substitutes the functional form for $G(n)$ into import demand (8):

$$\frac{M_{jt}^F}{M_{jt}} = s \cdot \bar{G} \cdot \left[1 - \exp(-\lambda n_{jt})\right] + u_{jt}. \qquad (11)$$

Because the model implies this relationship exactly, without an error term, we assume that u_{jt} is classical measurement error orthogonal to the decision to enter import markets and hence the number of imported inputs n.

Estimation. We now outline our empirical strategy and then explain the logic of identification. Technical details are relegated to Appendix B. Estimation proceeds in two steps. First we estimate (10) and (11) jointly using nonlinear least squares, which, given our structural assumptions, can also be thought of as a generalized method of moments estimation. To do this we nonparametrically estimate the unknown control function $h(\cdot)$ separately for each year as a third-order polynomial of z and investment, in which we also include a full set of 22 2-digit ISIC industry indicators.[17] Given the identifying assumptions that ε and u

---

[15]In contrast to our assumptions about $G(n)$, this functional form, which asymptotes to $\bar{G}$ as n grows without bound, has no maximum. Our implicit assumption, which is supported by our estimates, is that N_g is large enough that the difference between $\bar{G}$ and $G(N_g)$ is negligible. We obtained similar empirical results when we experimented with other functional forms for $G(n)$ including splines.

[16]Because z always includes calendar time, in effect we have a separate h function by year.

[17] In our main specifications we do not include nonlinear functions of the industry indicators because, in combination with the 12 year indicators, they would in effect create 264 groups—many of them small—for which a separate control function needs to be estimated, reducing power. Instead, below we report separate

are orthogonal to the right hand side variables, we obtain consistent estimates for β, γ, λ, δ and s. We then compute $a\bar{G} = \delta/\gamma$. To separately identify a and $\bar{G}$, we note that $1 - \bar{G}$ should equal the expenditure share of services in total spending on intermediates, which we compute directly from a manufacturing input-output table.[18] The second step is to estimate the coefficients of state variables z such as capital, which we do the same way as Olley and Pakes.

Identification. The difficulty in estimating the production function (9) predicted by the model is that ω is potentially correlated with all variables on the right hand side, including $G(n)$. For example, more productive firms tend to import more kinds of products. By estimating (10) instead of (9), we substitute out productivity using structural assumptions in the spirit of Olley and Pakes. The identification of the import demand equation also follows from structural assumptions, which restrict the functional form on the right-hand side of (11). To see how natural threats to identification are resolved, consider the concern that more productive firms both spend more on imports and import a greater number of varieties, which could introduce spurious correlation between the left hand side and $G(n)$ in (11). Importantly, our estimation is immune to this concern: productivity, which is explicitly incorporated in the model, cancels out of (11) because the left hand side is the *share* of imports in intermediate spending. While more productive firms do import more, they also spend more on intermediate goods, and given the homogenous production function, TFP drops out when we compute the ratio of these quantities. In fact, our structural assumptions yield a version of (11) which holds *exactly*, with no error term—this is why, given our model, u should be interpreted as classical measurement error.

It is useful to understand the variation which identifies our key parameters. Since we jointly estimate (10) and (11), the form of $G(n)$ is determined as the shape traced out by the import share and by output when n varies, controlling for firm productivity. Once we have this shape, a and s are estimated from the coefficients of $G(n)$, and are thus identified from variation in n given controls. In effect, we are comparing the output of two equally productive firms who import a different number of varieties. In the model, such variation in n comes from the fixed cost f, which affects the optimal number of imported inputs. This step requires that the distributions of f and F have enough variation so that n is not spanned by our set of controls, which clearly holds in the data.

---

estimates for several large industries, in which all parameters, including the control function, can differ across industries.

[18]When we later explore industry level estimates, we compute the service share separately by industry.

As with all structural estimation, the validity of our identification is guaranteed only if the model is correctly specified. One important possible misspecification is that firms might differ in their efficiency of import use A. Such variation can generate heterogeneity in s and a, which in turn can create correlation between u and $G(n)$. We incorporate such heterogeneity in our empirical strategy by allowing different groups of firms to have a different A.

Recovering deep parameters. Once a and s are estimated, we recover the deep parameters A and θ which govern the strength of the quality and imperfect substitution channels. The basic idea is that a high a combined with a low s shows the importance of imperfect substitution: even though imports are attractive (high a), the firm still uses a substantial share of domestic inputs (low s), indicating that mixing the two inputs is essential. Formally, we solve for A and θ from the following equations implied by the model:

$$-\ln(1-s) = a \cdot (\theta - 1), \tag{12}$$

$$\ln A = a \left[1 - \frac{\ln\ s}{\ln(1-\ s)}\right]. \tag{13}$$

The first equation states that import demand s is positively related to both the gains from importing a and the degree of substitution $\theta - 1$. Intuitively, a given gain from imports a maps into higher import demand when the elasticity is larger because importers are more willing to switch to foreign goods when they are good substitutes.[19] Because we have both s and a, this equation can be used to compute θ. The second equation (13) can then be used to back out the quality effect A. Here the intuition is that the difference between the quality effect $\ln A$ and the total gain from imports a is a reflection of imperfect substitution, which in turn is related to the import share s by the first equation.[20]

Consistency and standard errors. Our approach can be viewed as a GMM estimation and hence yields consistent estimates under our identifying assumptions. We obtain standard errors for all estimates from a bootstrap described in Appendix B.

Fixed costs. Finally we obtain measures, using some additional assumptions, for the fixed costs of importing. The idea is to invert equation (6), which expresses the optimal number of imported varieties n as a function of f and F. Implementing this idea raises two conceptual issues. First, a firm's choice of n depends not just on the import costs but also on the demand curve it faces. To make inference possible, we therefore assume that each

---

[19]This intuition is slightly imprecise because a reflects both the quality and the imperfect substitution effects.

[20]Our approach here parallels Feenstra (1994) and Broda and Weinstein (2006). They express the productivity (welfare) gain from variety as $x^{1/(1-\theta)}$, where $x < 1$ is the new expenditure share of old varieties. In our model, $x_i = 1 - s$ and the productivity gain is $\exp(a) = (1-s)^{1/(1-\theta)}$.

firm j faces an isoelastic demand curve $Y_{jt} = O_t p_{jt}^{-\eta}$, where we set $\eta = 5$, which is in the range of estimates reported by Broda and Weinstein (2006), and the demand shifter O_t is chosen so that we match total annual industry output in the data. Then equation (6) yields an explicit expression for n as a function of the fixed costs.

The second conceptual issue is that, because n is a discrete variable, it does not pin down the exact values of f and F. To deal with this problem, we first use (6) and the choice of n to compute, for each importer, upper and lower bounds for f and an upper bound for F. We then make additional assumptions to parametrically estimate the probability distribution of the costs. We assume that $\log f$ is normally distributed with constant variance and a mean which depends linearly on z_{jt} and ω_{jt}^{obs}, and we regress the log lower bound of f on these variables to obtain the coefficients.[21] We also assume that $\log F$ is normal with group-specific mean and variance, and compute these parameters by matching the share of importers as well as the mean difference in log output between importers and non-importers in each group.[22] Intuitively, higher fixed costs yield a lower share of importers, and higher dispersion in costs yields a smaller difference between the size of importers and non-importers.

4.2 Results

Table 3 summarizes our basic results. Because the dependent variable in the production function is log total sales, not value added, the coefficients of capital and labor are smaller than in the more common value-added specifications, while material costs have a large coefficient. All columns include year and 2-digit ISIC industry fixed effects (of which there are 22) so that our parameters are identified from within year and within industry variation across firms.

Column 1 reports the results from our empirical procedure in a baseline specification.[23] We estimate a highly significant per product import gain a of 0.174. This point estimate implies that the composite of the foreign and the domestic good is about $\exp(.174) - 1 = 19$ percent more efficient per dollar spent than the domestic good in itself. The share of non-

---

[21]Because the bounds we obtain are tight, estimating the regression using the upper bounds yields essentially identical results.

[22]Armenter and Koren (2009) calibrate the fixed costs of *exporting* the same way.

[23]Besides the year and industry effects, the control function $h(\cdot)$ used to take out unobserved productivity also contains third-order polynomials of investment and capital, estimated separately for each year and separately for domestic and foreign firms who might face different markets and hence a different investment decision. We also estimated specifications with separate control functions for the 22 2-digit industries (not reported) and obtained similar results.

Table 3: Baseline estimates

Dep. Var.: log sales	Baseline (1)	Conditioning on exporter (2)	Conditioning on past imports (3)	Nonlinear least squares (4)
Capital (α)	0.029 (0.003)	0.029 (0.003)	0.031 (0.002)	0.023 (0.002)
Labor (β)	0.200 (0.003)	0.200 (0.003)	0.199 (0.003)	0.202 (0.003)
Materials (γ)	0.788 (0.003)	0.784 (0.003)	0.782 (0.003)	0.791 (0.003)
Per-product import gain (a)	0.174 (0.046)	0.123 (0.038)	0.057 (0.03)	0.146 (0.037)
Import share (s)	0.666 (0.108)	0.674 (0.108)	0.674 (0.109)	0.670 (0.104)
Efficiency of imports (A)	1.116 (0.077)	1.083 (0.054)	1.037 (0.03)	1.098 (0.060)
Elasticity of substitution (θ)	7.301 [5.56;14.26]	10.071 [7.11;22.61]	20.791 [11.12;478.91]	8.571 [6.55;14.33]
Curvature of $G(n)$ (λ)	0.020 (0.007)	0.019 (0.007)	0.019 (0.007)	0.020 (0.007)
Foreign ownership	0.039 (0.011)	0.034 (0.011)	0.041 (0.011)	0.046 (0.008)
Exporter		0.044 (0.004)	0.045 (0.004)	
Previous importer			0.012 (0.005)	
Industry indicators	yes	yes	yes	yes
Year effects	yes	yes	yes	yes
P-value of test for A=1	0.008	0.008	0.008	0.004
Olley-Pakes control function				
Separately estimated:	By year	By year	By year	N/A
Includes third-order polynomials of	Investment, capital, foreign ownership	As column 1 and exporter status	As column 2 and past importer status	N/A
Observations	130,027	130,027	130,027	130,027

Notes: Bootstrapped standard errors clustered by firm are in parenthesis. For the elasticity of substitution (θ) we report a 95 percent confidence interval computed the same way in brackets. Columns 1-3 use the structural estimation procedure of Section 4.1 and column 4 is estimated from nonlinear least-squares. Industry indicators are 2-digit ISIC fixed effects. See Sections 4.1-4.2 and Appendix B for details.

service inputs among all intermediate inputs from the input-output table is $\bar{G} = 0.825$, and in column 1 the elasticity of output to intermediate inputs is estimated to be $\gamma = 0.788$. Combining these numbers, we predict that if a non-importer starts importing all tradeable varieties, it will experience an increase in log productivity of $a\bar{G}\gamma = 0.113$, or a productivity gain of about 12 percent.

The table also reports our estimates of the structural parameters A and θ. In the baseline specification, the price-adjusted quality of foreign products relative to their domestic counterparts is $A = 1.116$, which is significantly greater than one. Imported inputs are thus 12 percent better than domestic per dollar of expenditure, and this difference in price-adjusted quality accounts for about 61 percent of the per-product import gain. The remaining 39 percent comes from imperfect substitution: we find that the elasticity of substitution between domestic and foreign goods is $\theta = 7.3$. The basic empirical fact underlying the importance of imperfect substitution is that, in spite of the large gain from imports, the difference in the import share of firms who purchase more versus fewer foreign varieties is modest.

Column 2 re-estimates the model adding an indicator for export market participation in the production function. The reason is to distinguish the effect of imports from the "international engagement" of the firm, and to control for linkages between importing and exporting such as those emphasized by Kasahara and Lapham (2008). The control function in the Olley-Pakes procedure is now separately estimated by exporting status each year. The import estimates are slightly smaller but similar to the previous specification, suggesting that our procedure succeeds in isolating the impact of imports on productivity. Column 3 explores the possibility that entering import markets entails a sunk, rather than a fixed cost. When F is sunk, the past importing status of the firm becomes a state variable, which should therefore be included both in the regression and in the control function. The estimates are now smaller, but the per product import gain is still significantly different from zero, and the results are qualitatively similar to our baseline. Taken together, these results show that imports have a significant productivity effect across specifications.

A robust finding in columns 1-3 of the Table is that imperfect substitution is responsible for $35 - 40$ percent of the gains from importing. This result is consistent with the conclusions of Goldberg et al. (2009) who show, in micro data from India, that firms combine foreign and domestic varieties to increase their product scope: our results imply that combining these inputs also raises productivity. Our empirical finding that imperfect substitution amplifies the effect of higher quality inputs (i.e., that $a > \log A$) parallels theoretical arguments that complementarities between inputs can generate large income differences across countries. As Jones (2011) explains: "high productivity in a firm requires a high level of performance along

20

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 support for the productivity effect of this sort of interdependence in the context of combining foreign and domestic intermediate inputs.

Finally, column 4 reports the results from a specification which is the analogue of an OLS regression for our setting, in which equations (9) and (11) are jointly estimated with nonlinear least squares, but without using the Olley-Pakes procedure. Because this estimation does not fully control for firms' endogenous choice of variable inputs, we expect the coefficients of labor and material to be upward biased due to the standard reverse causality problem which plagues OLS estimates of productivity. In the table these coefficients are almost identical across columns, suggesting that the reverse causality bias is modest in practice. This result is not uncommon in the literature; one possible explanation in our setting is that the year and industry fixed effects pick up a substantial part of the endogenous variation.

Foreign versus domestic firms. Foreign firms, defined as companies in which foreigners have majority ownership, have played a very significant role in the Hungarian economy. In particular, in our data their sales share increased from 21 to 77 percent during 1992-2003. Table 3, in which all specifications include an indicator for foreign ownership, shows that foreign companies are on average about 4 percent more productive than domestic ones, suggesting that their growing participation has had significant aggregate productivity effects in Hungary. This observation raises the question of whether foreign firms are more productive in part because the they use imports more efficiently. Indeed, these firms may find it easier to locate low-cost input suppliers abroad, may have more extensive know-how about foreign goods, and may face lower transactions costs.

To explore whether foreign-owned companies are more efficient in their import use, Table 4 reports estimates in which the price-adjusted quality of imports A is allowed to be different across firms with different ownership status. Maintaining the assumption that firms use the same technology, we restrict the elasticity of substitution θ to be the same for the two groups. The regression results in the first specification show a large and significant difference in the efficiency of import use: we estimate $A = 1.18$ for foreign owned firms, while $A = 0.95$, not significantly different from one, for domestic firms. The large difference in A maps into a correspondingly large difference in the log per-product import gain: we obtain $a = 0.24$

Table 4: The gains from importing for foreign and domestic firms

Dep. Var.: log sales	Cross-section (1)		Event study of change in ownership (2)			
	Domestic	Foreign	Always domestic	Before switching	After switching	Always foreign
Capital (α)	0.03 (0.003)		0.027 (0.002)			
Labor (β)	0.2 (0.003)		0.203 (0.004)			
Materials (γ)	0.788 (0.003)		0.784 (0.003)			
Per-product import gain (a)	0.116 (0.024)	0.241 (0.059)	0.126 (0.035)	0.125 (0.051)	0.191 (0.059)	0.288 (0.074)
Import share (s)	0.435 (0.078)	0.695 (0.089)	0.444 (0.082)	0.442 (0.088)	0.59 (0.096)	0.739 (0.092)
Efficiency of imports (A)	0.949 (0.061)	1.182 (0.095)	0.953 (0.074)	0.952 (0.084)	1.081 (0.095)	1.25 (0.12)
Elasticity of substitution (θ)	5.936 [4.64;9.50]		5.672 [4.27;8.94]			
Curvature of $G(n)$ (λ)	0.024 (0.008)		0.022 (0.008)			
Foreign ownership	0.029 (0.013)		0.043 (0.012)			
Industry indicators	yes		yes			
Year effects	yes		yes			
P-value of test for A=1	0.398	0.004	0.552	0.418	0.299	0.004
P-value of test that A = previous column		0.004		0.972	0.099	0.016
Olley-Pakes control function						
Estimated separately:	By year		By year			
Includes third-order polynomials of	Investment, capital, foreign ownership		Investment, capital, foreign ownership			
Observations	109,555	20,472	102,661	2,048	3,792	21,525

Notes: Bootstrapped standard errors clustered by firm are in parenthesis. For the elasticity of substitution (θ) we report a 95 percent confidence interval in brackets. Both specifications are estimated with the structural procedure of Section 4.1. A different efficiency of imports (A) is estimated in specification (1) for foreign and domestic firms; and in specification (2) for always foreign firms, firms switching from domestic to foreign before the change, the same firms after the change, and always foreign firms. Industry indicators are 2-digit ISIC fixed effects. See Sections 4.1-4.2 and Appendix B for details.

for foreign and $a = 0.12$ for domestic companies. These results also imply that domestic companies benefit from imports primarily through imperfect substitution.

From a policy perspective it is important to understand whether foreign firms' greater efficiency in import use A is *caused* by them being foreign owned, or is due to other mechanisms such as selection, whereby foreign investors purchase firms which are better at using imports. To explore this question, we look for changes in the efficiency of import use in firms whose ownership status changes during our sample period. Intuitively, if foreign owners selectively purchase companies which are better at using imports, then the firm's A should remain the same after the company changes ownership. In contrast, if foreign owners improve import use because of lower transactions costs or other mechanisms, then we expect A to increase after the company changes owners. Plausibly, this adjustment might occur slowly, over the course of multiple years, because the company needs to learn the more efficient way of using imports.

In the second specification of Table 4 we estimate a separate price-adjusted import quality A for the following four groups of firms. (i) "Always domestic": firms who are not foreign-owned at any point in the panel. (ii) "Before switching:" firms who are domestically owned in the current year but whose ownership status will change at a later date in the sample. (iii) "After switching:" the same firms in years following the change in ownership. (iv) "Always foreign" firms who are foreign-owned throughout our data. The table shows that switchers have $A = 0.95$ before the change—essentially identical to the estimate of A for always domestic companies—which increases to $A = 1.08$ once they become foreign-owned. The values of A before versus after the change in ownership are significantly different at a p-value of 9.9 percent. We also find that the efficiency of import use for always-foreign companies is $A = 1.25$. Keeping in mind that we do not expect immediate full adjustment under either hypothesis, and that the small sample size (there are 938 switcher firms in the data) does not allow for a more precise estimate, we interpret our results from ownership changes as strong suggestive evidence that foreign firms' greater efficiency in import use is at least partly caused by them being foreign-owned. This evidence for causality suggests a potential policy complementarity between financial and trade liberalization, which we explore in greater detail in the next section.

Import effects by year and industry. To explore the robustness of our estimates and learn more about the impact of foreign goods on the Hungarian economy, we next explore how the gains from importing vary over time and across industries. Table 5 reports the efficiency of import use A estimated separately for foreign and domestic companies over three year intervals. Consistent with our earlier findings, throughout the sample period foreign firms

Table 5: The gains from importing over time

Dep. Var.: log sales	1992-94	1995-97	1998-00	2001-03
Efficiency of imports				
Domestic firms	0.928	0.984	0.950	0.962
	(0.071)	(0.072)	(0.078)	(0.09)
Foreign firms	1.052	1.320	1.198	1.151
	(0.115)	(0.164)	(0.092)	(0.107)
Other parameters (common across periods)				
Capital (α)		0.030		
		(0.002)		
Labor (β)		0.200		
		(0.003)		
Materials (γ)		0.788		
		(0.003)		
Elasticity of substitution (θ)		5.758		
		[4.242;8.511]		
Curvature of $G(n)$ (λ)		0.022		
		(0.007)		
Foreign ownership		0.029		
Industry indicators		(0.015)		
Year effects		yes		
P-value of domestic A=1	0.255	0.569	0.471	0.49
P-value of domestic A=previous column		0.451	0.706	0.804
P-value of foreign A=1	0.510	0.020	0.020	0.020
P-value of foreign A=previous column		0.059	0.392	0.725
Olley-Pakes control function				
Estimated separately:		By year		
Includes third-order polynomials of		Investment, capital, foreign ownership		
Observations		130,027		

Notes: Bootstrapped standard errors clustered by firm are in parenthesis. For the elasticity of substitution (θ) we report a 95 percent confidence interval in brackets. A different efficiency of imports (A) is estimated for foreign and domestic firms in each 3-year period using the structural procedure of Section 4.1. Industry indicators are 2-digit ISIC fixed effects. See Sections 4.1-4.2 and Appendix B for details.

Table 6: The gains from importing by industry

ISIC	Sector	Number of observations	Per-product import gain (a)	Import share (s)	Efficiency of imports (A)	Elasticity of substitution (θ)
15	Food and beverages	19,441	1.304***	0.935	3.568	3.100
18	Apparel	5,259	0.489**	0.938	1.613**	6.670
20	Wood products	8,061	0.228**	0.810	1.220	8.270
22	Printing and publishing	15,258	0.223***	0.697	1.168***	6.350
25	Rubber and plastics	7,534	0.194***	0.724	1.156***	7.650
26	Non-metallic minerals	5,252	0.05***	0.424	0.973	12.070
28	Fabricated metal products	18,155	0.237***	0.801	1.227**	7.790
29	Machinery	14,203	0.124***	0.465	0.973	6.070
33	Instruments	5,339	0.28**	0.796	1.271*	6.660

Notes: Table reports industry estimates of our baseline specification for industries with more than 5,000 observations. Significance levels (* at 10 percent, ** at 5 percent, *** at 1 percent) for the tests a=0 and A=1 are obtained from a bootstrap clustered by firm. In each industry, the Olley-Pakes control function is estimated separately by year and includes third-order polynomials of investment, capital and foreign status. See Sections 4.1-4.2 and Appendix B for details.

are better in using imports. The table also shows a large increase in the efficiency of import use between the 1992-94 and the 1995-97 period. This finding is consistent with the rapid trade liberalization taking place in Hungary in the early 1990s, which reduced the effective price of imports, increasing price-adjusted quality A. The 1992 Interim Agreement with the European Economic Community phased out most tariffs—the import-weighted tariff rate declined from 7.8 percent in 1992 to 3.2 percent in 1996—and Hungary joined the Central European Free Trade Agreement in 1992.

We also examine how the import effect varies by industry. Because different industries might face different production possibilities and a different market structure, we re-estimate our baseline specification separately for each of the 9 ISIC industries in which there are more than 5,000 firm-year observations. We allow for different capital, labor and material coefficients as well as a different $G(n)$ and different Olley-Pakes proxy functions for each industry. Table 6 reports our estimates of the key model parameters by industry.[24] Because each industry has a smaller number of observations, the estimates are noisier, but the table confirms the main patterns identified earlier. Imports have a significantly positive productivity effect in all 9 industries; and imperfect substitution is responsible for 42 percent of

---

[24]Significance levels are indicated by stars for the per-product import gain a and for the efficiency of imports A.

these gains on average. The table also shows some interesting variation across industries. The largest per-product import gain—over 130 percent—is in food and beverages, while the smallest—about 5 percent—is in non-metallic minerals. Our estimates imply that the difference between these two industries is due to the combination of the two productivity channels: imports are imperfectly substitutable and high-quality in the former industry, while they are close substitutes and relatively low quality in the latter. It seems also intuitively plausible that inputs in non-metallic minerals should be more substitutable then those in the food industry. The industry results also highlight how deep parameters are determined by our coefficient estimates. For example, both the wood products industry and printing and publishing have a per-product import gain of about 22 percent; but because in wood products the import share estimate is larger, our model implies a higher elasticity of substitution. Intuitively, given the quality advantage of foreign goods, a higher import share must come from greater substitutability.

Fixed costs. Table 7 reports summary statistics for the estimated fixed cost distribution both for the baseline specification of column 1 in Table 3, and for the first specification in Table 4 which distinguishes the efficiency of import use for foreign and domestic firms. In all cases, we find that the median costs are low but dispersion across firms is high. For example, in the baseline case while the typical firm pays an entry cost F of only \$143, the mean entry cost is a much higher \$1,301, reflecting the variance. Intuitively, our model explains the fact that many mid-sized and large firms do not import by assigning a high fixed cost to them, pulling up both the variance and the average of F. When we distinguish firms by ownership, we find that foreign firms find it cheaper to start importing: they typical foreign firm needs to pay only \$39 in contrast to the typical domestic firm which pays \$158. For foreign firms, surprisingly, we also estimate a higher median fixed cost f (\$892 versus \$331). Although, as we have shown in Fact 1, foreign firms import more kinds of products even conditional on their size, our estimates imply that this is because they gain more from imports, *not* because of lower average fixed costs. Intuitively, given that they find imports very productive (high A), the only reason in our model why some foreign firms do not import a large number of varieties is that they face a high per product cost.

Finally, we measure the effect on our productivity estimates of the fixed costs of importing. As we discussed in Section 3.2, $a\gamma G(n)$ is the proper measure of the productivity gain from importing only when fixed costs are ignored. To measure the quantitative importance of fixed import costs, we compute, in the baseline specification, the average of $F + nf$ among all importers, and obtain \$10,173. In contrast, the average cost saving from importing among all importers in our sample is \$611,493. Thus the fixed costs account for

Table 7: Fixed costs

		Entry cost (F)			Per-product cost (f)		
		median	mean	log stdev	median	mean	log stdev
Baseline		$143	$1,301	2.17	$443	$657	0.785
By ownership status	Domestic	$158	$1,107	1.97	$331	$451	0.718
	Foreign	$39	$671	2.38	$892	$1,617	1.082

Notes: Summary statistics implied by our estimates for the fixed costs of importing are reported in 1991 USD. Baseline refers to specification (1) in Table 3 and By ownership status refers to specification (1) of Table 4. See Sections 4.1-4.2 and Appendix B for details.

only about 1.7 percent of the total cost savings created by imports, and as a result, ignoring them does not significantly alter the aggregate implications of our model. The reason for the small contribution of fixed costs is that much of the cost-saving from importing is realized by mid-sized and large firms importing multiple varieties. For these firms the fixed cost of importing the first few—most important—products is negligible relative to the large savings obtained through the reduction in marginal cost.

5   Applications

This section develops two applications of our estimates. In Section 5.1 we quantify the aggregate productivity effects of imports in Hungary, and in Section 5.2 we explore the implications of tariff policies in our estimated economy.

5.1   Decomposing the productivity gains in Hungary

To decompose the aggregate productivity gain in Hungary into various channels, we first write a firm's residual log productivity $\omega_{jt} = 1_{jt}^{F}\mu + \rho_{jt}$. Here 1_{jt}^{F} is an indicator for foreign status, μ is the Hicks-neutral mean log productivity premium of foreign firms, and ρ_{jt} measures remaining variation in productivity. Then the (log) productivity of firm j in year t is

$$\phi_{jt} = [1_{jt}^{F} \cdot a_{t}^{F} + 1_{jt}^{D} \cdot a_{t}^{D}] \cdot \gamma G(n_{jt}) + 1_{jt}^{F}\mu + \rho_{jt}$$

where a_{t}^{F} and a_{t}^{D} denote the per product import gain for foreign and domestic firms in year t and $1_{jt}^{D} = 1 - 1_{jt}^{F}$ is an indicator for domestic firms. Following Olley and Pakes (1996), we

measure aggregate TFP as the sales-weighted average of firms' log TFP

$$\Phi_t = \sum_i \sigma_{jt}\phi_{jt}, \tag{14}$$

where σ_{jt} is the output share of firm j in year t. Denoting by $\bar{G}_t^D$ and $\bar{G}_t^F$ the sales-weighted average of $G(n_{jt})$ and by σ_t^D and σ_t^F the sales share for domestic and foreign firms in year t, simple algebra shows that the growth in aggregate productivity between time t and time 0 equals

$$\begin{aligned}
\Phi_t - \Phi_0 = {}& [(a_t^D - a_0^D)\gamma\sigma_0^D\bar{G}_0^D + (a_t^F - a_0^F)\gamma\sigma_0^F\bar{G}_0^F] \\
& + [a_t^D\gamma\sigma_0^D(\bar{G}_t^D - \bar{G}_0^D) + a_t^F\gamma\sigma_0^F(\bar{G}_t^F - \bar{G}_0^F)] \\
& + [a_t^D\gamma(\sigma_t^D - \sigma_0^D)\bar{G}_t^D + a_t^F\gamma(\sigma_t^F - \sigma_0^F)\bar{G}_t^F] \\
& + \mu(\sigma_t^F - \sigma_0^F) + \left[\sum_j (\sigma_{jt}\omega_{jt} - \sigma_{j0}\omega_{j0})\right].
\end{aligned}$$

The five terms on the right-hand side have the following intuitive interpretation. The first two terms measure the gains from growth in imports, created by an increase in the price-adjusted quality of imports (first term) and an increase in the number of varieties imported (second term). The third term reflects the greater aggregate effectiveness of using imports in the economy, due to the increase in the share of foreign firms. Taken together, these terms represent the productivity gains created by importing. The fourth term measures the direct effect of increased foreign presence, caused by the higher average residual productivity of foreign-owned firms μ. The final term measures additional, non-import-related variation in productivity.

We use the above expression to decompose productivity growth in Hungary. To smooth out business-cycle fluctuations and ensure a large number of observations, we use all firm-year observations in the three-year range 1992-1994 as our "starting date" and similarly all observations during 2001-03 as our "ending date," but interpret the results as a decomposition for 1993-2002. We use the coefficient estimates from Table 5 which is estimated in our main sample, but, because our goal is to compute trends over time, we compute the decomposition in the firm-level sample defined in Section 2.1. Specifically, we calculate ω_{jt} for each firm-year (of the firm-level sample) during 1992-94 as well as during 2001-03 from the production function (9) and the coefficients in Table 5. Similarly, we calculate the various sales shares and averages of $G(n)$ in the firm-level sample for all observations in 1992-94 as well as 2001-03. The results are summarized in Table 8. Our numbers imply that total productivity growth in the manufacturing sector in Hungary during this period was 15 percent.

Table 8: Productivity growth in the Hungarian manufacturing sector 1993-2002

Growth in aggregate productivity	15.0%
intensive margin of imports	0.9%
extensive margin of imports	2.2%
increased import by foreign firms	2.9%
direct effect of foreign firms	1.4%
other	7.6%

Notes: Total productivity growth between the periods 1992-94 and 2001-03, interpreted as growth from 1993 to 2002, is decomposed into five different margins as described in section 5.1.

More than one-third of this growth, about 6 percent, can be attributed to various import-related mechanisms. Increased imports generated a productivity gain of 3.1 percent, most of which, 2.2 percent, comes from more firms importing more kinds of products. This result confirms the quantitative significance of the "new goods margin" also emphasized in the context of product scope by Goldberg et al. (2009). The more efficient use of imports by an increasingly foreign-owned manufacturing sector adds another 2.9 percent, highlighting the substantial aggregate effect of the interaction between foreign capital and importing. These large numbers indicate that imports were a significant contributor to economic growth in Hungary.[25]

5.2 Tariff effects

Motivated by the large aggregate effects of importing, we now turn to explore how economic policies affect import-related productivity gains. Our main focus is on the effect of tariff cuts. Intuitively, a reduction in tariffs can increase productivity by reducing the price of imports and thereby increasing price-adjusted quality A, and also by increasing the number of varieties the firm chooses to import. Our goal is to understand how the magnitude of these gains depends on the economic environment and on other concurrent economic policies.

Model economy. We simulate tariff cuts in a static partial equilibrium economy constructed based on the data and our estimates. Firms behave as in our theoretical model, with technology parameters (in most experiments) given by the coefficients in the first spec-

---

[25]The remaining 9 percent of productivity growth is due to factors unrelated to importing. Higher foreign presence, consistent with the empirical results of Brown, Earle and Telegdy (2006), had a substantial direct effect by virtue of foreign firms being more productive; and we also find a 7.6 percent increase caused by forces outside our analysis.

ification of Table 4. In particular, foreign firms differ from domestic firms in that they use imports more effectively and that they have, on average, higher residual log productivity. The population of firms in the model economy is the union of all firm-year observations in our main sample. To model heterogeneity, we exogenously assign to each firm its actual capital stock and foreign status from the data, as well as its ω^{obs} and ε as predicted by the regression coefficients.

In the absence of point estimates, we randomly draw the fixed importing costs f and F, separately for foreign and domestic firms, from their estimated fixed cost distributions (Table 7). Random assignment ensures that in the model economy the fixed costs of non-importers and importers are drawn from the same distribution. To solve for their optimal choices, we also assume—as in the estimation of fixed costs—that in the market for their final good firms face an isoelastic demand curve with price elasticity $\eta = 5$. All firms face the same input prices, which we set exogenously to match aggregate output and input usage, and then we let firms make their optimal decisions about labor, materials, imports, and the price of their final good. As we assumed when estimating the model, firms can only respond to the observed component of productivity $\omega^{obs} = h(I, z)$, and not to ε which affects only realized output but none of the inputs.

In the counterfactual experiments we track aggregate outcomes as firms' optimal choices change in response to changes in prices, fixed costs, ownership status, and other parameters. For simplicity and to highlight intuitions transparently, our experiments ignore general equilibrium price effects.[26] We also hold capital fixed to avoid dynamic considerations. Endogenizing capital would likely amplify the effects we document for output and aggregate productivity due to reallocation of capital to its most productive use.

Tariff policy. We first study the effects of a uniform input tariff change of size τ, which changes the price of all foreign inputs by a factor of $(1+\tau)$ relative to the benchmark economy just defined. For example, $\tau = 0$ is our benchmark model, while $\tau = 0.1$ corresponds to a tariff increase of ten percentage points. In our partial equilibrium setting, the tariff change affects the economy by altering the price-adjusted import quality A of all firms.

Figure 1 plots changes in aggregate (log) productivity as a function of the tariff change in our simulated economy in three scenarios, which differ in the share of foreign firms in the economy. In the benchmark scenario the same firms are foreign in our model as in the data, while in the other two scenarios all firms are domestic respectively all firms are foreign. In keeping with the construction of the model economy, when we change the ownership status

---

[26] The partial equilibrium approach can also be justified if the policies we evaluate affect a small industry in a small open economy.
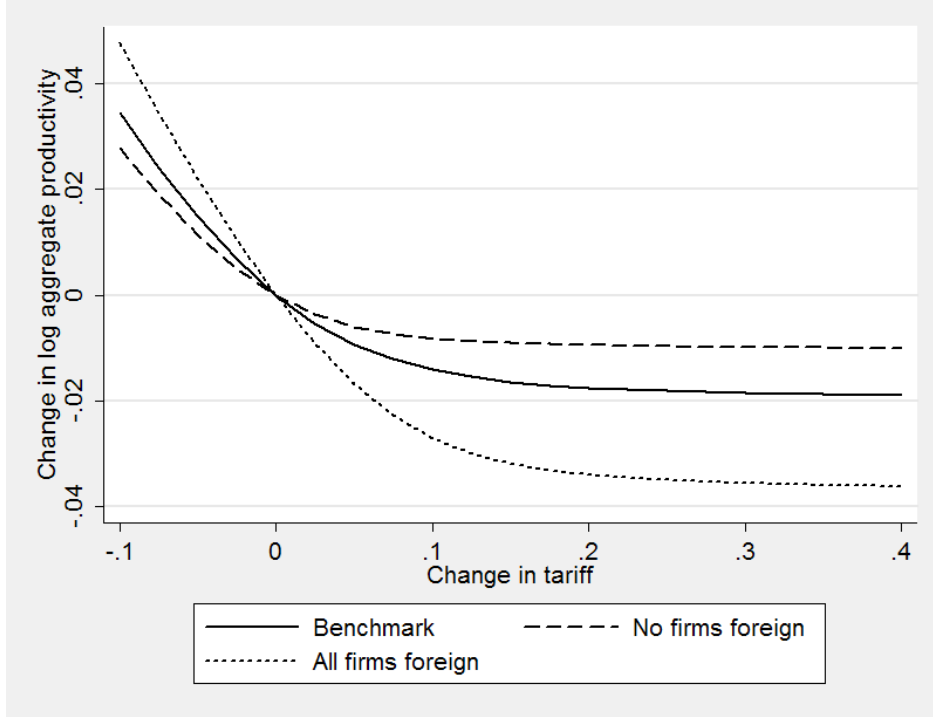
Figure 1: Effect of tariff changes on aggregate productivity

of a firm, we change its efficiency of using imports A, redraw its fixed costs f and F from the appropriate distribution, and adjust its Hicks neutral log productivity ω by the mean difference between firms in the two groups. We then compute the sales-weighted aggregate log productivity defined in (14) for each value of τ, and—to measure the productivity *change* in response to the tariff policy—normalize this quantity separately in each scenario (by adding a constant) to zero at the point of no tariff change.

Two key lessons emerge from Figure 1. First, the fact that the curves are convex implies that the gains in aggregate log productivity from a given tariff cut are larger when the initial tariff is low, i.e., in a more open economy. Such "increasing returns" to tariff cuts may perhaps seem surprising, but the underlying intuition is straightforward. A marginal reduction in tariffs increases productivity by reducing the cost of foreign inputs; and this cost reduction is higher when more firms use more kinds of foreign inputs. Importantly, convexity also implies that larger cuts have a more-than proportional effect on log productivity, because they also increase the set of imported goods on which the associated cost-savings occur. The second lesson of the Figure is that the gains from a tariff cut are larger in an economy with greater foreign presence. This follows because foreign firms are more effective in using imports.

Table 9: Counterfactual experiments

Panel A			
Tariff reduction	High fixed cost	Baseline	Low fixed cost
40% to 30%	0.00%	0.00%	0.30%
10% to 0%	0.10%	1.40%	3.50%

Panel B			
Tariff reduction	No firms foreign	Baseline	All firms foreign
40% to 30%	0.00%	0.00%	0.10%
10% to 0%	0.80%	1.40%	2.70%

Notes: Table reports changes in aggregate TFP in our simulated economy in response to a 10 percentage point tariff reduction under various scenarios. High fixed costs are 3 times the baseline and low fixed costs are 1/3 of the baseline for each firm in the simulated economy. See Section 5.2 for details.

To quantify the two forces identified by the Figure and explore their broader policy implications, Table 9 computes the change in aggregate productivity that results from a 10 percentage point reduction in tariffs in several hypothetical scenarios. As above, a non-zero tariff τ alters the values of A for both foreign and domestic firms relative to the benchmark economy. Extending the intuition about the interaction between openness and tariff cuts, Panel A focuses on the combination of decreasing tariffs and changing the *fixed costs* of importing, such as those associated with import licensing. The three columns correspond to environments which differ in the fixed costs of importing. For each firm, we assign f and F which are three times their baseline values for that firm in the high fixed costs column, equal to the baseline values in the middle column, and one-third of the baseline values in the low fixed costs column.[27]

Parallelling the larger productivity gain from liberalization when tariffs are lower, Table 9 shows a larger productivity gain when the fixed costs are lower. For example, in the first row, a reduction in tariff from 40 percent to 30 percent has essentially no productivity effect when fixed costs are at their baseline level because almost no firms are importing, but does increase productivity by 0.3 percent in the low-fixed-cost environment. Moreover, as the second row shows, this complementarity between fixed costs and tariff cuts is stronger when the economy is more open to start with: a tariff cut from 10 percent to zero increases

---

[27]In the absence of direct evidence on how liberalization affects fixed costs, our scenarios are motivated by broad patterns in the World Bank's *Doing Business* survey. In the average OECD country, it takes 14 days to start a new business and 11 days to import a standard containerized cargo. These time costs are about three times as high, 45 days and 38 days, respectively, in Sub-Saharan Africa.

productivity by 1.4 percent in the baseline, and by a much larger 3.5 percent in the low-fixed-cost environment. Taken together, these results highlight a policy complementarity between reducing tariffs and reducing the fixed costs of importing, and also show that this complementarity is stronger in a more open economy with lower initial tariffs.

Panel B of Table 9 documents results about a similar policy complementarity between tariff cuts and *FDI liberalization*, particularly in open economies. The first row shows that a tariff reduction from 40 percent to 30 percent has essentially no effect both in the case when the economy has no foreign firms and when the composition of foreign and domestic companies is as in our baseline. But, as the second row shows, the same 10 percentage point tariff cut in a more open economy generates productivity gains of 0.8 percent with no foreign firms and 1.4 percent when the share of foreign firms is as in our baseline. Intuitively, in a fairly closed economy the 10 percentage point tariff cut does not make imports cheap enough for either foreign or domestic firms, and is hence ineffective. But in a more open economy, the same tariff cut reduces the cost of inputs that firms actually use, and hence increases productivity. This effect is stronger in the presence of foreign firms, who are more productive in using imports.

Our predictions about policy complementarities, which we obtained using only Hungarian data, seem broadly consistent with the liberalization experience of the 1990s in India. Consistent with the first complementarity, tariff cuts in India, which were accompanied by dismantling substantial non-tariff barriers, lead to rapid growth in new imported varieties (Goldberg et al. 2009) and a large increase in firm productivity (Topalova and Khandelwal 2011). And consistent with the second complementarity, these effects were stronger in industries with higher FDI liberalization (Topalova and Khandelwal 2011).

Import substitution. We next explore the effect of tariffs on the demand for domestic intermediate goods. Our goal here is to contrast the implications of the quality and imperfect substitution mechanisms, and hence for simplicity we perform the analysis in a model where technological parameters are given by the estimates in column 1 of Table 3, so that foreign and domestic firms use imports equally efficiently. We analyze tariff effects in the following three scenarios. (1) Foreign and domestic goods are perfect substitutes: the benefit of importing is entirely due to price-adjusted quality. (2) Foreign and domestic goods have the same price-adjusted quality: the benefit of imports is entirely due to imperfect substitution. (3) About 39 percent of the gains are caused by imperfect substitution, as in our baseline estimate. We implement these scenarios by holding fixed the per-product import gain $a = 0.174$ as in our estimate, and by adjusting A and θ for the different scenarios. For instance, in the first scenario we set $\theta = 100$ which implies effectively perfect substitution, and let

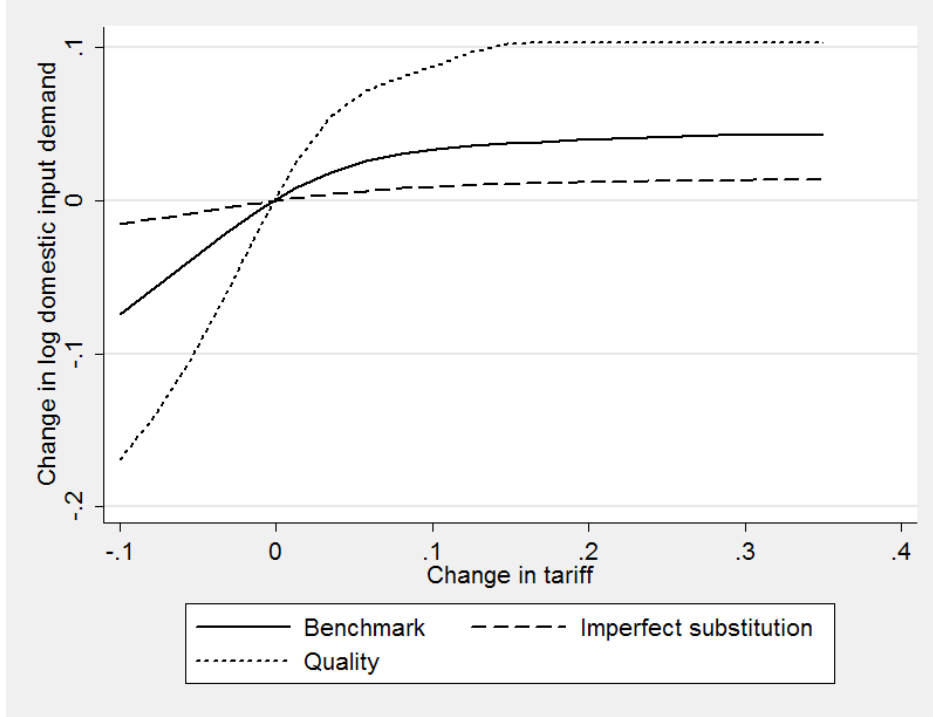Figure 2: Effect of tariff changes on domestic input demand

$A = \exp(a)$.

Figure 2 plots the change in the log dollar value of domestic input demand relative to the baseline in response to a change in tariffs. Domestic input demand is more sensitive to tariffs when the benefit of imports comes from quality than when it comes from imperfect substitution. This is intuitive: when foreign goods are close to perfect substitutes, even small price changes can bring about large import substitution. Because our estimates assign a large role to imperfect substitution, the middle curve, corresponding to the empirically estimated composition of the two channels, is fairly flat. This curve also reflects the additional affect that the losses caused by the tariff cut are counteracted by increased demand for all inputs created by the productivity gains from importing. One implication of these two forces is that redistributive losses due to import substitution may not be very large in practice. 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.

6 Conclusion

This paper explored the effect of imports on productivity by estimating a structural model of importers in a panel of Hungarian firms. Perhaps the main caveat to our analysis is that, in the absence of plausibly exogenous variation, identification requires us to use with full force the restrictions imposed by our structural framework. This caveat may be weighted against the benefit that a formal model allows for counterfactual policy analysis. A potentially interesting research area is to seek reduced-form evidence for our new predictions, such as the complementarities between different liberalization policies, or the effects of tariffs on domestic input demand.

A second important domain that our analysis neglects is the use of capital goods. Work by Caselli and Wilson (2004) suggests that capital imports are potentially significant because of the technology embedded in them. Careful analysis of the productivity gains from capital imports would further contribute to our understanding of the link between international trade and economic growth.

References

Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction, *Econometrica* **60**(2): 323–351.

Amiti, M. and Konings, J. (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from Indonesia, *American Economic Review* **97**(5): 1611–1638.

Armenter, R. and Koren, M. (2009). Economies of scale and the size of exporters, Working paper. FRB Philadelphia and Central European University.

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).

Bernard, A. B., Jensen, J. B., Redding, S. J. and Schott, P. K. (2007). Firms in international trade, *Journal of Economic Perspectives* **21**: 105–130.

Broda, C. and Weinstein, D. E. (2006). Globalization and the gains from variety, *The Quarterly Journal of Economics* **121**(2): 541–585.

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

Brown, J. D., Earle, J. S. and Telegdy, A. (2006). The productivity effects of privatization: Longitudinal estimates from Hungary, Romania, Russia, and Ukraine, *Journal of Political Economy* **114**(1): 61–99.

Caselli, F. and Wilson, D. J. (2004). Importing technology, *Journal of Monetary Economics* **51**(1): 1–32.

Das, S., Roberts, M. J. and Tybout, J. R. (2007). Market entry costs, producer heterogeneity, and export dynamics, *Econometrica* **75**: 837–873.

Ethier, W. J. (1979). Internationally decreasing costs and world trade, *Journal of International Economics* **9**: 1–24.

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**(1): 157–77.

Goldberg, P., Khandelwal, A., Pavcnik, N. and Topalova, P. (2009). Imported intermediate inputs and domestic product growth: Evidence from India, *Quarterly Journal of Economics* **125**(4): 1727–1767.

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

Grossman, G. M. and Rossi-Hansberg, E. (2008). Trading tasks: A simple theory of offshoring, *American Economic Review* **98**(5): 1978–97.

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

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators, *Econometrica* **50**(4): 1029–54.

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. (2011). Intermediate goods and weak links in the theory of economic development, *American Economic Journal: Macroeconomics* **3**(2): 1–28.

Kasahara, H. and Lapham, B. (2008). Productivity and the decision to import and export: Theory and evidence, Working paper, University of Western Ontario and Queens University.

Kasahara, H. and Rodrigue, J. (2008). Does the use of imported intermediates increase productivity? Plant-level evidence, *Journal of Development Economics* **87**(1): 106–118.

Kremer, M. (1993). The O-ring theory of economic development, *The Quarterly Journal of Economics* **108**(3): 551–75.

Markusen, J. R. (1989). Trade in producer services and in other specialized intermediate inputs, *American Economic Review* **79**: 85–95.

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.

Schor, A. (2004). Heterogeneous productivity response to tariff reduction. Evidence from Brazilian manufacturing firms, *Journal of Development Economics* **75**(2): 373–396.

Topalova, P. and Khandelwal, A. (2011). Trade liberalization and firm productivity: The case of India, *Review of Economics and Statistics* **93**(3): 995–1009.

Tybout, J. R. (2003). Plant- and firm-level evidence on "new" trade theories. In Harrigan and Choi (eds): Handbook of International Economics, Blackwell, Malden.

Appendix

A   Proofs

Proof of equation (9). Given the Cobb-Douglas structure, total expenditure on intermediates M must equal their price index times their Cobb-Douglas aggregate:

$$M = \prod_{i=1}^{N} P_i^{\gamma_i/\gamma} \prod_{i=1}^{N} X_i^{\gamma_i/\gamma}. \tag{15}$$

Taking logs and combining this equation with (3) with (1) yields

$$y = \alpha k + \beta l + \gamma m + a \sum_{i=1}^{n} \gamma_i + \omega \tag{16}$$

as desired. We can also compute the log of the total cost of the firm as

$$\ln C = \frac{1}{\alpha + \beta + \gamma} \left[(1 - \alpha - \beta - \gamma)y + \alpha r + \beta w - \omega - a \sum_{i=1}^{n} \gamma_i \right]. \tag{17}$$

Thus $\phi = a \sum_{i=1}^{n} \gamma_i + \omega$ drives both the variation in output and the variation in costs across firms.

Profits when demand is isoelastic. We write demand for the firm's product as $Y = OP^{-\eta}$ where O is a demand shifter and P is the price of the final good, or, with lowercase variables denoting logs,

$$y = o - \eta p.$$

The firm charges a constant markup, hence the price is proportional to marginal cost. From the production function, and holding capital fixed, the (log) short-run marginal cost is

$$c_0 + \frac{1 - \beta - \gamma}{\beta + \gamma} y - \frac{1}{\beta + \gamma} [\omega + \alpha k + a\gamma G(n)] + \frac{\beta}{\beta + \gamma} w + \frac{\gamma}{\beta + \gamma} p,$$

where c_0 is a constant. It is easy to see that the ratio of profits to revenue is a constant $\pi_0 = 1 - (\beta + \gamma)(1 - 1/\eta)$. In turn log revenue equals

$$p + y = c(o, p, w) + \frac{\eta - 1}{\beta + \gamma} [\omega + \alpha k + a\gamma G(n)], \tag{18}$$

where $c(o, p, w)$ captures the effect of changes in prices and demand. Letting $\phi = (\eta - 1)/(\beta + \gamma)$, we can therefore write the gross profit of the firm, ignoring fixed costs, as

$$\pi^G(k, \omega, n, c) = \pi_0 \exp \left\{ c + \phi \left[\omega + \alpha k + a\gamma G(n) \right] \right\} \tag{19}$$

where $c = c(o, p, w)$ is exogenous to the firm. We can use this profit function to (i) estimate fixed costs consistent with observed choices of n, and (ii) to determine counterfactual choices of n. The profit function is increasing in n, at a decreasing rate for the baseline estimated parametrization for $G(n)$.

B Estimation and counterfactuals

B.1 Estimating the coefficients

First step. For any given value of λ, the shape of the $G(n)$ function is pinned down, and hence our two equations (10) and (11) become linear and can be separately estimated. Therefore we first fix λ and estimate the equations separately, and then search for λ to obtain the best fit.

We estimate (11) using OLS for a given λ, in which, to ensure that the estimates match the total volume of import flows, we weight each observation jt by total intermediate expenditure M_{jt}. We estimate (10) using the unknown h proxy function as follows. We regress all left-

and right-hand-side variables of (10) on third-order polynomials of I_{jt} and z_{jt} (with the exception that industry indicators only enter linearly) to obtain a conditional expectation $E(x_{jt}|I_{jt}, z_{jt})$. Subtracting these conditional expectations from (10) and making use of the fact that $E(h(I_{jt}, z_{jt})|I_{jt}, z_{jt}) = h(I_{jt}, z_{jt})$, we can write the production function as

$$y_{jt} - E(y_{jt}|I_{jt}, z_{jt}) = \beta[l_{jt} - E(l_{jt}|I_{jt}, z_{jt})] + \gamma[m_{jt} - E(m_{jt}|I_{jt}, z_{jt})] + \delta\{G(n_{jt}) - E[G(n_{jt})|I_{jt}, z_{jt}]\} + \varepsilon_{jt}.$$

This equation can be estimated by OLS because ε_{jt} is orthogonal to all choice variables. We then search for the $\hat{\lambda}$ that minimizes the sum of squared residuals, $\sum_{jt} \varepsilon_{jt}^2 + u_{jt}^2$ via Newton's method.

Second step. The second step follows Olley and Pakes (1996) exactly. The coefficients of all state variables (capital, foreign ownership, exporter and past importer status) are estimated in this step. We first regress exit on a third-order polynomial of investment and z_{jt} using the same algorithm as in the first stage. The predicted exit, $\hat{p}_{it}$, together with past productivity $\omega_{i,t-1}$ can then provide sufficient statistics about the bias in ω_{it},

$$E(\omega_{jt}|\text{exit}_{j,t+1} = 0, \omega_{j,t-1}) = \psi(\hat{p}_{jt}, \omega_{j,t-1}),$$

where ψ is an unknown function. Current productivity is a nonparametric function of past productivity and the bias coming from endogenous exit. We approximate this function by third-order polynomials.

We can rewrite the conditional expectation above as a regression equation,

$$\omega_{jt} = \psi(\hat{p}_{it}, \omega_{j,t-1}) + e_{jt}$$

where e_{jt} is a TFP innovation uncorrelated with past productivity and exit. As we do not observe past productivity, we measure it as past output minus the combination of inputs,

$$\hat{\omega}_{j,t-1}(\zeta) \equiv y_{j,t-1} - \hat{\beta}l_{j,t-1} - \hat{\gamma}m_{j,t-1} - \hat{a}\hat{\gamma}G(n_{j,t-1}) - \zeta z_{j,t-1},$$

where ζ is the coefficient vector of the state variables, in particular, it includes the capital coefficient α. We then estimate ζ by nonlinear least squares from

$$\omega_{jt} = \sum_{i=0}^{2} \sum_{m=0}^{2-i} \hat{\omega}_{j,t-1}(\zeta)^i \hat{p}_{jt}^m + e_{jt}.$$

This estimate is consistent because e_{jt} is orthogonal to all right-hand-side variables.

Because our estimation procedure can be viewed as a generalized method of moments, it is consistent by Hansen (1982).

B.2 Bootstrap to obtain standard errors and tests

We obtain standard errors, confidence intervals and p-values from a bootstrap with 500 repetitions. We sample firms with replacement, holding their entire time path together so as to preserve the joint distribution of variables at different points in time. This is akin to clustering standard errors by firm.

Standard errors are obtained as the empirical standard deviation of the 500 estimates. For the elasticity of substitution θ we report 95 percent confidence intervals. This is the range between the 2.5th and 97.5th percentile of the empirical distribution of $\hat{\theta}$.

We perform two-tailed tests for the hypothesis $A = 1$ as follows. If the estimate $\hat{A} > 1$, we count the number of realizations in the bootstrap for which $\hat{A}_{(i)} \leq 1$ and report the p-value of $2 \cdot n/500$. Similarly, if $\hat{A} < 1$, the p-value is based on the number of estimates for which $\hat{A}_{(i)} \geq 1$.

B.3 Estimating fixed costs

We set the elasticity of demand $\eta = 5$. To work with the profit function (19), we also need to specify $c(o, p, w)$, which is exogenous to the firm. We do this separately in each year to ensure that our model matches, given our coefficient estimates and the demand elasticity, total sales in the data. For any given value of c, for all firms in a given year we compute their optimal sales using (18) given their observed k, ω and n, and then sum across firms. We then search for the value of c which equates this measure of total sales with what we observe in the data.

Given c, we estimate the fixed costs as follows. If we observe a firm importing n_{jt} products in year t, then

$$\pi^G(k_{jt}, \omega_{jt}, n_{jt}, c) \geq \pi^G(k_{jt}, \hat{\omega}_{jt}, n_{jt} - 1, c) + f$$

and

$$\pi^G(k_{jt}, \hat{\omega}_{jt}, n_{jt}, c) \geq \pi_t(k_{jt}, \hat{\omega}_{jt}, n_{jt} + 1) - f.$$

These inequalities yield lower and upper bounds for f computable given our estimates and parameter choices which we denote $\underline{f}(k_{jt}, \hat{\omega}_{jt}, n_{jt})$ and $\bar{f}(k_{jt}, \hat{\omega}_{jt}, n_{jt})$.

We assume that actual fixed costs follow a lognormal distribution conditional on state variables. We let the mean of $\log f$ depend linearly on capital, productivity, and foreign ownership status as well as industry and year indicators: $\mu + bk_{jt} + c\omega_{jt}^{obs} + d1_{jt}^F + $ fixed effects. We estimate μ, b, c, d and σ from a regression of $\log \underline{f}$ on these variables in the sample of firms with $n_{jt} \geq 1$ for whom the upper and lower bounds are close to one another.

We assume that, for all firms within a group, F has log mean μ_F and log standard deviation σ_F. We calibrate these parameters by matching two moments of the data: the fraction of importing firms and the relative size (difference in log output) of importers and non-importers. A higher average fixed cost lowers the fraction of firms who find it profitable to enter the import market, and this monotonic relationship makes μ_F easy to calibrate. To calibrate dispersion σ_F, we note that if σ_F is small, firms have very similar fixed costs, and hence firm size is the main determinant of whether firms import. As a result, the output of importers and non-importers will be very different. In contrast, if σ_F is large, the decision to import will depend mostly on the realization of F, and hence the output of importers and non-importers will be similar.

B.4 Contribution of fixed costs to productivity

To compute an upper bound for the average fixed cost paid by importers in our data, we first calculate using our baseline specification, for each importer, the upper bound on the per-product fixed cost $\bar{f}(k_{jt}, \hat{\omega}_{jt}, n_{jt})$ as in Section B.3 and multiply it with the number of imported products n_{jt}. The average of these upper bounds across importers is \$9,987. To this we add the average entry cost F across importers in the simulated economy. We use the simulated economy because we only estimate the distribution of F but not its realizations. Although the unconditional mean of F is \$1,301, the majority firms with a high realization do not import, leading to a mean F across importers of only \$186. Overall average fixed costs of importing are therefore \$10,173.

To compute the cost savings from importing, we compute in the baseline specification of our simulated economy that if the average importer stopped importing, optimal production of the same amount of output would require a \$611,493 higher cost. From these numbers, fixed costs are only 1.7 percent of the cost saving due to imports.