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 and CEPR

March 2015

Abstract

We estimate a model of importers in Hungarian micro data and conduct counterfactual policy analysis to investigate the effect of imported inputs on productivity. We find that importing all input varieties used in production would increase a firm's revenue productivity by 22 percent, about half of which is due to imperfect substitution between foreign and domestic inputs. Foreign firms use imports more effectively and pay lower fixed import costs. Our estimates imply that during 1993-2002, a quarter of the productivity growth in Hungary was due to imported inputs. Simulations show that the productivity gain from a tariff cut is larger when the economy has many importers and many foreign firms, implying policy complementarities between tariff cuts, dismantling non-tariff barriers, and FDI liberalization.

*E-mail addresses: halpern@econ.core.hu, korenm@ceu.edu and szeidla@ceu.edu. We thank Márta Bisztray, István Ilyés and Péter Tóth for excellent research assistance, Pol Antràs, Péter Benczúr, Christian Broda, Jan De Loecker, Gita Gopinath, Penny Goldberg (the editor), Elhanan Helpman, Marc Melitz, Ariel Pakes, Roberto Rigobon, John Romalis, David Weinstein, two anonymous referees, and seminar participants for helpful comments. For financial support, we thank the Global Development Network (Award RRC IV-061 to Halpern and Koren) the Hungarian Scientific Research Fund (Award T048444 to Halpern and Koren) the Alfred P. Sloan Foundation (Szeidl) and the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) ERC grant agreements number 313164 (Koren) and 283484 (Szeidl).

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] A 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, 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 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. Firms must pay a fixed cost each period for each variety they choose to import. Imported inputs affect firm productivity through two distinct channels: as in quality-ladder models they may have a higher price-adjusted quality, and as in product-variety models they imperfectly substitute domestic inputs.[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 capital, labor, materials, and a 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

[1] Results are conflicting for Brazil: Schor (2004) estimates a positive effect while Muendler (2004) finds no effect of imported inputs on productivity. And for Argentina Gopinath and Neiman (2013) show that variation in imported inputs may have contributed to fluctuations in aggregate productivity.

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

component of productivity. We also build on the approach of De Loecker (2011) to control for demand effects, and follow Gandhi, Navarro and Rivers (2013) in estimating the materials coefficient from input demand. Given these controls, the import effect is 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 and face the same level of demand, 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.

Our results show that the productivity gains from imported inputs are substantial. In the baseline specification, increasing the fraction of tradeable goods imported by a firm from zero to 100 percent would increase revenue productivity by 22 percent and quantity productivity by 24 percent. We continue to estimate large productivity gains from importing when—as in models in which the cost of importing is sunk, rather than fixed—measures of the firm's past importing behavior are included as state variables. These results suggest that imported inputs play a significant role in shaping firm performance in the Hungarian economy.

We then turn to decompose the import effect into the quality and imperfect substitution channels. We first note that for a given productivity gain from importing a 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 which differ in the number of imported varieties. We find that combining imperfectly substitutable foreign and domestic varieties is responsible for about half 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. We say that a firm "has been foreign owned" if either on the current date or on some past date its majority owners were foreigners.[4] Because they have know-how about foreign

[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.

[4]The vast majority of firms that had been foreign owned at some past date remained foreign owned for the duration of their life in our sample. Our definition reflects our view that foreign ownership has lasting effects on firm operations.

markets and can access cheap suppliers abroad, these firms may gain more from spending on imports. This is an important possibility because firms that had been foreign owned played a central role in Hungary: during 1992-2003, their sales share in manufacturing increased from 21 percent to 80 percent. When we re-estimate our model allowing for differences in the efficiency of import use by ownership status, we find that firms that have been foreign owned benefit about 24 percent more than purely domestic firms from each dollar they spend on imports. We also conduct an event study of ownership changes which yields suggestive evidence that part of the premium in the efficiency of import use is caused by foreign ownership. This result implies a potential complementarity between foreign presence and importing.

Our analysis also yields estimates of the product-level fixed costs of importing. We find that—as in the model of Gopinath and Neiman (2013)—these costs increase in the number of imported products, and also that the fixed cost schedule of firms that have been foreign owned is below that of domestic firms. Lower import costs are thus a second factor generating higher benefits from importing to foreign firms.

In the second half of the paper, we develop two applications to study the economic and policy implications of our estimates. We first quantify the contribution of imports to productivity growth in Hungary during 1993-2002. Our estimates imply a productivity gain of 21.1 percent in the Hungarian manufacturing sector, of which 5.9 percentage points, more than one quarter, can be attributed to import-related mechanisms. Approximately 80 percent of these import-related gains are due to the increased volume and number of imported inputs, while the other 20 percent is the result of increased foreign ownership in combination with foreign firms being better at using imports. Thus imports contributed substantially to economic growth in Hungary, and the complementarity between foreign presence and importing had a sizeable aggregate effect. These results complement the findings of Gopinath and Neiman (2013) who emphasize the role of imported inputs in driving *fluctuations* in aggregate productivity.

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 foreign inputs, should raise both firm-level and aggregate productivity. Our main result is that the size of the aggregate productivity gain depends positively on two broad features of the environment: (1) the initial import participation of producers; (2) the degree of foreign presence. Perhaps surprisingly, higher initial import participation—either due to low tariffs or to low fixed costs—implies larger gains from a tariff cut. This is because the set of inputs whose prices are affected is larger, and hence firms save more with the tariff cut. 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. Because foreign firms are more effective in using imports, a similar complementarity exists between tariff cuts and FDI liberalization. These complementarities seem broadly consistent with the liberalization experience in the early 1990s in India. Consistent with the fixed cost 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 foreign ownership complementarity, these effects were stronger in industries with higher FDI liberalization (Topalova and Khandelwal 2011).

Our tariff experiment also highlights the differential implications for domestic input demand of the quality and imperfect substitution mechanisms. When the benefit of imports comes from quality differences, domestic import use—in an intermediate range—is quite sensitive to tariffs. In contrast, when the benefit from imports comes from imperfect substitution, domestic input use is a relatively flat function of tariffs. This difference is intuitive: when foreign goods are close to perfect substitutes, even a small price change can bring about large import substitution. Another force is that losses to domestic input suppliers caused by a tariff cut are partially offset by increased demand for their products due to increased productivity.[5] Because our estimates assign a significant role to imperfect substitution, and because of the second force, we obtain a relatively inelastic demand curve for domestic inputs. One lesson from this analysis is that the magnitude of redistributive losses due to import substitution depend strongly on the extent of substitution and on the initial level of tariffs. More broadly, identifying the specific mechanism driving the effect of trade policies can help evaluate the impact of these policies in other dimensions.

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) and Bernard, Jensen, Redding and Schott (2012). 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, Kasahara and Lapham (2008) who investigate the link between exports and imports, and De Loecker, Goldberg,

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

Khandelwal and Pavcnik (2014) who study the effect of trade liberalization on markups. 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 Section 5 describes the results. In Section 6 we use the estimates to conduct counterfactual analysis. We discuss some caveats with our approach in the concluding Section 7.

2 Data

2.1 Data and sample definition

Main data sources. Our panel of essentially all Hungarian manufacturing firms during 1992-2003 is created by merging balance sheet data and trade data for these firms. Firms' balance sheets and profit and loss statements come from the Hungarian Tax Authority for 1992-1999, and from the Hungarian Statistical 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, with the main omitted category firms being individual entrepreneurs without employees. The data for 2000-2003 include all firms with at least 20 employees and a random sample of firms with 5-20 employees. We thus lose some firms in 2000. These firms, 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.

Data on firms' annual export and import value, disaggregated by products at the 6-digit Harmonized System (HS) level, come from the Hungarian Customs Statistics. Because the 6-digit classification is noisy, we aggregate the data to the 4-digit level. In the rest of the paper we use the terms "product" and "good" to refer to a HS4 category.[6] Because we are interested in the effect of imported inputs, we use data on those imported products which are classified as intermediate goods, industrial supplies or capital good parts in the Broad Economic Categories classification. We merge the balance sheet and trade data using unique numerical firm identifiers.

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

While we have product level data on *imported* input purchases, a limitation is that—because balance sheets only measure total spending on intermediate goods—we do not have corresponding product-level data for *domestic* input purchases. 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 firms' 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. In our data the total value of intermediate imports by wholesalers and retailers is about 2 percent of total intermediate input use by all firms in all sectors. This fact suggests that in our data the role of intermediation for inputs is relatively small, 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 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. However, because these goods cross the border, processing trade is recorded in our trade data. This inconsistency creates problems: in several observations, the value of *imported* intermediate inputs, as measured by customs, exceeds the value of *all* intermediate inputs, as measured by the balance sheet. Similarly, some firms' exports in the customs data are substantially higher than their exports in the balance sheet data.

To deal with this reporting problem, we construct a measure of each firm's processing trade. This measure is defined as the difference, when it is positive, between customs exports and balance sheet exports. We classify a firm as a "processer" in a given year if the ratio of processing trade to balance sheet sales exceeds 2.5 percent. This cutoff is approximately the median across observations in which the ratio is positive. With this definition, about 9 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.

Sample definitions. 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.[7] 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. Because it has more observations, unless otherwise noted we will use the main sample in our analysis. After all exclusions, 127,472 firm-year observations remain in this sample.

Variable definitions. For each firm in each year, the balance sheet data contain information on the ownership shares of domestic and foreign owners. We say that firm j in year t "has been foreign owned" if either in that year or in some prior year foreigners had majority ownership. This definition is motivated by the view that foreign ownership has lasting effects on a firm's operations. It also solves the problem that for some firms ownership data is missing in some years. Reflecting the fact that only a quarter of the 5,009 firms that have been foreign owned ever switch back to majority domestic ownership, we sometimes simply refer to a firm which has been foreign owned as "foreign."

Because firms must file balance sheets in the county in which they are headquartered, we can classify each firm in each year as being located in one of the 20 counties in Hungary (19 actual counties and the city of Budapest). The firm-level data also contain information on the firm's industry. We work with the 2-digit International Standard Industrial Classification (ISIC, revision 3) industry definitions, and for firms that report different industries in different years, we assign the most common industry reported.

Other data sources. We obtain 2-digit industry-level input and output price indices for 1992-2003 from the Hungarian Statistical Office. We also exploit an industry-level input-output table which was constructed for the year 2000 by the Hungarian Statistical 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.

Fact 1. There is substantial heterogeneity in the import patterns of firms. Half of firms do not import at all; firms which are larger or have been foreign owned are more likely to import.

This fact can be seen by comparing across columns in Table 1. This table presents summary statistics for several key variables in our main sample separately for importing and

[7] 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	48.85	17.42	98.87
Sales (thousand USD)	2,870	420	6,770
Capital per worker (thousand USD)	13.75	11.22	17.76
Sales per worker (thousand USD)	51.32	36.99	74.11
Material share in output	0.82	0.79	0.87
Exporter indicator	0.35	0.15	0.68
Export share in output	0.10	0.04	0.20
Importer indicator	0.39	0.00	1.00
Import share in materials	0.10	0.00	0.27
Number of imported products (HS4)	4.71	0.00	12.22
Foreign owned	0.19	0.09	0.34
State owned	0.03	0.02	0.04
Number of observations	127,472	78,273	49,199
Number of firms	26,593	20,921	13,341

Notes: Table 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 1998 dollars.

non-importing firms. Importers employ about 6 times as many workers and sell about 16 times as much as non-importers. Importers are also more frequently foreign and more likely to export.[8]

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 an indicator for whether the firm has been foreign owned shows that doubling firm size is associated with a 25 percent increase in the number of imported products, and, conditional on size, firms which have been foreign owned import 187 percent more products than purely domestic firms.

The patterns shown here are 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.

[8]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.

Table 2: Import dynamics

Year	Total imports (USD millions)	Import growth (%)	Contribution to import growth: Intensive margin (pp)	New firms (pp)	New importers (pp)	New products (pp)	Stopping firms (pp)	Stopping importers (pp)	Dropped products (pp)
1992	1,968								
1993	2,358	19.8	16.2	9.6	5.1	10.6	-12.5	-2.6	-6.5
1994	3,182	35.0	28.2	3.0	5.6	8.1	-3.2	-1.3	-5.5
1995	3,885	22.1	18.6	3.3	1.9	5.6	-1.3	-1.6	-4.3
1996	5,194	33.7	12.0	2.1	4.6	24.2	-3.2	-1.0	-4.9
1997	6,668	28.4	26.1	2.0	1.2	3.9	-1.8	-0.3	-2.8
1998	8,875	33.1	32.1	0.7	0.6	3.0	-0.5	-0.3	-2.4
1999	10,799	21.7	20.4	1.4	0.4	2.2	-0.8	-0.3	-1.7
2000	13,421	24.3	21.7	8.0	0.6	1.8	-6.5	-0.1	-1.3
2001	16,413	22.3	21.3	0.2	1.6	2.8	-2.0	-0.5	-0.9
2002	15,425	-6.0	-3.6	0.4	0.4	1.2	-3.1	-0.4	-0.9
2003	14,521	-5.9	-4.5	0.4	1.0	1.1	-1.3	-0.1	-2.5
Average		20.8	17.1	2.8	2.1	5.9	-3.3	-0.8	-3.1
1992 to 2003		637.8	111.4	513.4	11.5	53.1	-38.2	-0.8	-12.5

Notes: Total imports are in 1998 dollars. The contributions to import growth columns measure the percentage point increase in imports attributable to different mechanisms and sum to the import growth column. The intensive margin measures (net) growth in imports of products that the firm also imported the previous period (the previous year, and in the last row 1992). New firms are firms that did not exist in the previous period, new importers are firms that existed but did not import in the previous period, new products are newly imported products of existing importers. Stopping firms, stopping importers and dropped products are defined analogously.

To document this fact, for each firm, we order imported products by their share in the total import spending of the firm. 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 3.4 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. The table decomposes the growth in imported intermediate inputs in the manufacturing sector into a within-firm intensive margin and six different extensive margins: new firms, new importers, new imported products; and exiting firms, firms stopping to import, and within-firm shedding of imported products. It is instructive to look at the average of these decompositions over all years, reported in the second to last row. On average, im-

ports of intermediate inputs grew by 20.8 percent per year. This growth can be decomposed into a within-firm intensive margin, which contributed 17.1 percentage points; growth on the three extensive margins (firms, importers, products) which contributed 10.8 percentage points; and decline on the three extensive margins which contributed -7.1 percentage points. Among the extensive margins, firms adding new imported products was the biggest contributor (5.9 percentage points). The large magnitude of the extensive margin calls for an explicit model of the decision to enter additional import markets. And the comparable magnitudes of the margins associated with adding and shedding imported inputs (5.9 and 3.1 percentage points) suggest that the decision to import likely entails some per-period fixed costs.[9]

The last row in the table reports a similar decomposition for the entire 1992-2003 period. During this time imports grew by about 693 percent. The main component of this growth, explaining 571 percentage points, is the "new firms" margin: imports by firms that did not exist in 1992. This fact suggests that manufacturing in Hungary underwent substantial restructuring during our sample period. One of our goals in this paper is to examine the productivity implications of this restructuring and the associated increase in importing.

3 An Industry Equilibrium Model of Imported Inputs

Motivated by the above stylized facts, in this section we build a static model of industry equilibrium in which firms use both domestic and imported intermediate goods for production.

3.1 Setup

Production technology. Firms in industry s are indexed by $j = 1, ..., J^s$. The output of firm j is given by the production function

$$Q_j = \Omega_j K_j^{\alpha} L_j^{\beta} \prod_{i=1}^{N} X_{ji}^{\gamma_i}, \tag{1}$$

where K_j and L_j denote capital and labor used in production, X_{ji} denotes the quantity of intermediate composite good i used by firm j, and Ω_j is Hicks neutral total factor produc-

[9]Due to the change in sample definition, we lose some importing firms in 2000 (see Section 2.1). These observations are classified as exiting firms, but 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.

tivity (TFP). The Cobb-Douglas weight γ_i measures the importance of intermediate input i for production. Motivated by Fact 2, we allow γ_i to be different for different goods i. The total weight of all intermediate goods is $\gamma = \sum_i \gamma_i$. We assume that the production structure—characterized by the parameters α, β, γ_i, and the set of intermediate inputs—is the same for all firms in industry s.

Each intermediate good X_{ji} is assembled from a combination of a foreign and a domestic variety:

$$X_{ji} = \left[(B_{ji} X_{jiF})^{\frac{\theta-1}{\theta}} + X_{jiH}^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}, \tag{2}$$

where X_{jiF} and X_{jiH} 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 we assume that the firms are price takers in these input markets. The price-adjusted quality advantage of the foreign input is $A_{ji} = B_{ji} P_{iH} / P_{iF}$. Intuitively, A_{ji} measures the advantage of a dollar spent on a foreign relative to a domestic variety.

We make several simplifying assumptions about intermediate inputs. To allow for non-tradeable inputs in a simple way, we assume that they coincide with the set of services, and assign an infinitely high foreign price and hence $A_{ji} = 0$ to them. We can then estimate the input share of non-tradeables from an input-output table. We also assume that the price-adjusted quality A_{ji} of all tradeable goods used by firm j is the same across inputs within a group of firms: $A_{ji} = A$. This assumption simplifies our analysis and still allows us to estimate the *average* quality advantage of imports. Note, we do not restrict $A > 1$, because we also want to allow foreign goods to have potentially lower quality than domestic goods. When estimating the model, in some specifications we permit A to depend on characteristics such as year or whether the firm has been foreign owned. 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 fixed costs to access foreign intermediate inputs. Similarly to Gopinath and Neiman (2013) we assume that firm j faces a fixed cost schedule: when it is already importing $i-1$ intermediate inputs, importing an additional input requires an incremental fixed cost $f_j^i \geq 0$. Thus if firm j imports i types goods, it pays a total cost of $f_j^1 + f_j^2 + ... + f_j^i$. We denote $\bar{f}_j = (f_j^1, ..., f_j^{N_g})$. To make the model consistent with the high frequency of exit from import markets, when estimating the model we assume that these costs are due every period.

Uncertainty. We assume that the log of Ω_j can be written as $\omega_j = \omega_j^{obs} + \varepsilon_j$ where the firm observes ω_j^{obs} before it makes import choices, but it observes ε_j only after all choices have been made.

Demand. Demand for goods in industry s is determined by the preferences

$$U(\{Q_j\}_{j=1}^{J^s}) = \left[\sum_{j=1}^{J^s} V_j^{1/\eta} Q_j^{(\eta-1)/\eta}\right]^{\eta/(\eta-1)} \tag{3}$$

where η is the elasticity of substitution between products and V_j is a demand shifter associated with the product of firm j. We normalize $\sum_{j=1}^{J_s} V_j = 1$. To ensure that a solution to the firm's profit-maximization problem exists, we also assume $\alpha + \beta + \gamma < \eta/(\eta - 1)$.

Timing. We assume that K_j and L_j are predetermined, and use the model to understand how input purchases, output, revenue and price are determined in equilibrium.

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 firm's efficiency advantage (per dollar of spending) when using a foreign, rather than a domestic, input. 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. 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 that involve product-level information only for imported, but not domestic input purchases.

3.2 Model solution

Input choices. We first consider the gain from importing a particular intermediate input i. The effective price of the composite good X_{ji} if the firm chooses to import variety i can be found by solving the cost-minimization problem associated with (2):

$$P_{ji} = \left[P_{iH}^{1-\theta} + (P_{iF}/B_{ji})^{1-\theta}\right]^{1/(1-\theta)} = P_{iH}\left[1 + A^{\theta-1}\right]^{1/(1-\theta)} \tag{4}$$

using the notation that $A_{ji} = B_{ji}P_{iH}/P_{iF}$ and our assumption that $A_{ji} = A$. Because the price of the composite good X_{ji} is $P_{ji} = P_{iH}$ if the firm only uses the domestic input, the (log) percentage reduction in the cost of the tradeable composite good i when imports are

also used is

$$a = \frac{\log\left[1 + A^{\theta-1}\right]}{\theta - 1}. \tag{5}$$

Parameter a measures the *per-product import gain* and hence is of central interest to us. This parameter incorporates the cost-savings created by both the quality and the imperfect-substitution channels, and hence it is higher when the price-adjusted quality A is higher or when the degree of substitution θ is lower. Because of imperfect substitution, for finite θ the firm uses both domestic and foreign inputs, so that the optimal expenditure share of the foreign good in the total spending for variety i,

$$S = A^{\theta-1}/(1 + A^{\theta-1}) \tag{6}$$

satisfies $0 < S < 1$.

Connecting n_j to import demand and output. In choosing which varieties to import, the firm trades off the saving in marginal cost from using imports against the fixed cost of importing. Since the fixed cost schedule only depends on the number of imported products, and since the per-product gain a is the same for all products, a firm which imports n products will choose to import those with the highest γ weight, i.e., products $i = 1, ..., n$. We now use this observation to characterize how n_j affects import demand and output.

The following function measures the relative importance for production of the inputs the firm chooses to import:

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

Since $\gamma_1 \geq \gamma_2, ... \geq \gamma_{N_g} \geq 0$, 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 inputs in all intermediate inputs.

Now consider import demand conditional on n_j. Denoting expenditure on *all* intermediate inputs by $M_j = \sum_{i=1}^{N} P_{ji} X_{ji}$ and expenditure on *foreign* intermediate inputs by $M_j^F = \sum_{i=1}^{N} P_{iF} X_{jiF}$, the spending share on imports—a measure of import demand—equals

$$\frac{M_j^F}{M_j} = S\frac{\sum_{i=1}^{n_j} \gamma_i}{\gamma} = SG(n_j) \tag{8}$$

where S, defined in (6), is the optimal expenditure share of imports within a composite good. Intuitively, firms that import a greater number of products n_j have a larger share of foreign goods in total intermediate spending.

Next consider output conditional on n_j. Let $\varrho = \Gamma \prod_{i=1}^{N} P_{iH}^{\gamma_i/\gamma}$ denote the price of the composite of domestic intermediate inputs in industry s, where $\Gamma = \prod_{i=1}^{N} (\gamma_i/\gamma)^{-\gamma_i/\gamma}$ is a

constant. We assume that this is the price index the statistical office computes for industry inputs. For a firm which chooses to import n_j varieties and optimally chooses the composition of domestic and foreign inputs within each such variety, we show in Appendix A that the production function (1) implies

$$q_j = \alpha k_j + \beta l_j + \gamma(m_j - \rho) + a\gamma G(n_j) + \omega_j \tag{9}$$

where the lowercase variables q_j, k_j, l_j, m_j, ω_j denote logs and $\rho = \log(\varrho)$. The first three terms on the right-hand side measure the contribution to output of capital, labor, and intermediate inputs; the final term is the Hicks-neutral productivity shifter ω. The novelty in the equation is the fourth term, which represents the contribution of imports. Intuitively, a firm which chooses to import n_j varieties will have a percentage cost reduction of a on the associated composite inputs, the total weight of which is $\sum_{i=1}^{n_j} \gamma_i$. This cost reduction maps into a corresponding increase in output for a given total spending on intermediate inputs.

Industry equilibrium. To determine revenue and profits, we need to combine equation (9) with the demand for the firm's product. Let the industry output price index P be defined by $P^{1-\eta} = \sum_{j \in s} V_j P_j^{1-\eta}$, and let industry output be $Q = U(\{Q_j\}_{j=1}^s)$ as given by equation (3). Then, following De Loecker (2011), denoting $R_j = P_j Q_j$ and lowercase variables with logs, we can derive from (9) that

$$r_j - p = \frac{1}{\eta} q + \frac{1}{\eta} v_j + \alpha^* k_j + \beta^* l_j + \gamma^*(m_j - \rho) + \gamma^* a G(n_j) + \omega_j^*, \tag{10}$$

where star indicates that the coefficient is multiplied by $(\eta - 1)/\eta$, for example, $\alpha^* = \alpha(\eta - 1)/\eta$. The term on the left hand side is firm revenue normalized by the industry price index. The first two terms on the right hand side come from the demand system and correct for the fact that we express revenue rather than quantity. The remaining terms on the right hand side have similar interpretation as in (9), the difference being that they are now adjusted by the factor $(\eta - 1)/\eta$ to account for price effects.

Choosing the number of imported varieties. We now return to the choice of n_j. Let $\pi(n)$ denote expected operating profits (without subtracting the fixed costs of importing) if the firm imports n goods. Here the expectation is over the only source of residual uncertainty ε_j. Because of the constant elasticity of demand, expected operating profits are a constant fraction of expected revenue, and can be computed from (10).[10] The optimal import decision of the firm is then

$$n_j = \arg\max_{\tilde{n}} \pi(\tilde{n}) - \sum_{i=1}^{\tilde{n}} f_j^i. \tag{11}$$

[10]For notational simplicity we suppress the dependence of π on other firm-level variables such as k or ω^{obs}. We compute the profit function explicitly in Appendix A.

Imports augmenting productivity. It is natural to interpret equation (9) as a production function for output in which the firm's total factor productivity is given by $\phi_j = a\gamma G(n_j)+\omega_j$, 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 only when the fixed costs are small relative to the overall productivity gain. Because importing reduces *marginal* costs but requires the payment of *fixed* costs, this is more likely to hold for medium and large firms which import multiple different products.[11] In the empirical analysis, we will show that—because the bulk of production and importing is performed by mid-sized and large importers—on average in our data fixed costs are small relative to the cost-savings generated by imports. Hence in practice little is lost by treating ϕ as a measure of productivity, which is what we do below.

By a similar logic, it is natural to interpret (10) as a production function for revenue. In this expression revenue productivity—defined as revenue minus the contributions of capital, labor and intermediate inputs—equals $\phi_j^R = \frac{1}{\eta}q + \frac{1}{\eta}v_j + a\gamma^* G(n_j) + \omega_j^*$. Here the first two terms represent demand effects that influence revenue conditional on the contributions of the factors of production. As with quantity productivity above, here too little is lost by ignoring the role of fixed costs.

4 Estimation

4.1 Assumptions

We now state assumptions about dynamics and heterogeneity which allow us to estimate our static model in panel data. Consider a firm j in industry s, located in county c, in year t. Recall that $\omega_{jt} = \omega_{jt}^{obs} + \varepsilon_{jt}$, where ω_{jt}^{obs} is observable to the firm at the beginning of period t. Following De Loecker (2011) we also assume that the (log) within-industry demand shifter of firm j at time t can be written as $v_{jt} = \psi_0 + \psi \cdot d_{jt}$ where d_{jt} is an observable demand shifter.

[11] 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.

Building on Olley and Pakes (1996) we assume that conditional on a vector of state variables, the firm's investment decision is a monotone function of observed productivity ω_{jt}^{obs}. Formally, assume that $I_{jt} = \xi(\omega_{jt}^{obs}, k_{jt}, l_{jt}, z_{jt})$ where ξ is increasing in its first argument. It is natural that investment should depend on capital and labor, which are by assumption predetermined. We also allow I_{jt} to depend on a vector of state variables $z_{jt} = (d_{jt}, q_t^s, s, t, c, o_{jt})$. Here d_{jt} is the within-industry demand shifter and q_t^s measures industry-level demand. Both of these, as shown in Section 3.2, affect the firm's problem in period t. Because demand or productivity might evolve differently by industry, year and location, we also include s, t and c in z_{jt}. Finally o_{jt} denotes other potential state variables which might also affect the firm's investment decision (for example, through differential access to finance). We always include in o_{jt} an indicator for whether the firm has been foreign owned.

The timing for firm j within period t is the following.

1. Observe ω_{jt}^{obs}.
2. Observe the vector of state variables z_{jt}, decide whether to exit.
3. Decide on investment I_{jt}.
4. Observe the fixed costs of importing $\bar{f}_{jt}$, the wage w_t^s and the input price index ρ_t^s.
5. Decide on the number of imported products n_{jt} and total material spending m_{jt}.
6. Observe ε_{jt}.
7. Produce output q_{jt} and sell at a price determined by the demand curve.
8. Set $l_{j,t+1}$.

We assume that the productivity shocks ε_{jt} are i.i.d. and independent of all other shocks; that the fixed cost realizations $\bar{f}_{jt}$ are i.i.d. and independent of all other shocks conditional on z_{jt}; and that the industry-level factor prices (w_t^s, ρ_t^s) are independent of all other shocks and i.i.d. between industries and over time. Thus, consistent with the assumption that it determines firm investment, the vector $(\omega_{jt}^{obs}, k_{jt}, l_{jt}, z_{jt})$ fully characterizes the distribution of shocks facing firm j in period t.

We also assume that the observed component of productivity can be written as $\omega_{jt}^{obs} = \mu(s, c, o) + \varpi_{jt}$. Here the mean shifter $\mu(s, c, o) = \mu_s^1 + \mu_c^2 + \mu_o \cdot o$ so that the mean of ω^{obs} can, through fixed effects, vary by industry and by county, and can also depend linearly on the state variables in o_{jt}. And ϖ_{jt} is a Markov process satisfying $\varpi_{jt} = f(\varpi_{j,t-1}) + e_{jt}$

where e_{jt} are i.i.d. and independent of all other shocks. Finally, we require that for all firms the process $(\omega_{jt}^{obs}, z_{jt})$ is Markov with the same dynamics. It follows—again consistent with our assumption on the investment function—that the current realization of $(\omega_{jt}^{obs}, z_{jt})$ fully determines its distribution in future periods.[12] Assumptions about the dynamics of shocks similar to ours are frequently used in the productivity literature.

Our key variable of interest is the benefit of importing, measured by a. In the estimation we assume that observations can be partitioned into groups based on z_{jt}—for example by ownership status or year—such that the quality advantage of the foreign input A is constant for observations within a group, but may vary across groups. An implication is that the per-product import gain a, and also the import share measure S, will stay constant within, but vary across groups. We let $g = 1, ..., \bar{g}$ index groups.

Heterogeneity. Our framework allows for considerable heterogeneity. Firms can differ in their productivity, factor use, foreign and domestic intermediate input use, and also in their realized fixed costs. Crucially, we also permit heterogeneity across inputs through the γ_i parameters. We do assume that the γ_i—essentially, the $G(\cdot)$ function— are the same across firms. This assumption implies that additional varieties decline in importance identically across companies, but it *does not* imply that firms in different industries 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 industries.

4.2 Estimating the import effect in a single group

We begin by describing our estimation strategy for the case in which all firms have the same efficiency of import use A (that is, $\bar{g} = 1$). We will later discuss how to extend the procedure when there are multiple groups with different values of A.

We estimate our model using three equations. We use the empirical counterpart of the import share equation (8) to estimate the $G(n)$ function. We assume that $G(n)$ has the parametric functional form

$$G(n) = \begin{cases} \bar{G}\left(1 - \left[1 - \left(\frac{n}{\bar{n}}\right)^{\lambda}\right]^{1/\lambda}\right) & \text{if } n \leq \bar{n}, \\ \bar{G} & \text{if } n > \bar{n}. \end{cases} \tag{12}$$

[12] De Loecker (2011) allows ω_{jt}^{obs} to also depend on $d_{j,t-1}$, but does not include q_t^s in the vector of state variables. In contrast, while we do not permit d_{jt} to directly affect productivity, we do allow for persistence in the dynamics of d_{jt} and q_t^s and hence include both of them in z_{jt}. These variables are also included in the equation determining investment.

Here $\lambda \in (0,1)$ and $\bar{G} \in (0,1)$. This functional form yields a declining marginal benefit of additional imports, which eventually—when $n > \bar{n}$—completely levels off. The import share equation (8) yields our first estimating equation

$$\frac{M_{jt}^F}{M_{jt}} = S \cdot G(n_{jt}) + u_{jt} \tag{13}$$

where $G(n)$ is assumed to be given by the above function. Because the model implies this relationship exactly, without an error term, we assume that u_{jt} is measurement error orthogonal to the number of imported inputs n_{jt}. We use this equation to estimate the shape parameter λ and the import share coefficient S.

Our second estimating equation exploits the firm's fist order condition for intermediate inputs to connect the coefficient γ with the material share in production. Our use of this equation parallels the empirical approach of Gandhi et al. (2013). Because materials are chosen after all shocks except for ε_{jt}^* are realized, profit maximization and the Cobb-Douglas production function imply

$$\gamma^* \frac{E_\varepsilon(R_{jt})}{M_{jt}} = 1 \tag{14}$$

where E_ε refers to expectations taken with respect to the uncertainty in ε_{jt}. We use this equation to estimate γ^*.

Our third estimating equation comes from the revenue production function. Here the classic identification problem is that firm productivity ω_{jt}^{obs} can be correlated with other variables on the right hand side. We follow the Olley and Pakes (1996) approach in getting around this problem by inverting the monotone increasing investment function ξ to get

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

with an unknown h "control" function. Substituting this expression into (10), denoting $\delta^* = \gamma^* a$, and using $v_{jt} = \psi_0 + \psi \cdot d_{jt}$, we obtain our empirical import-augmented production function

$$r_{jt} - p_t^s = \frac{1}{\eta} q_t^s + \frac{1}{\eta}\psi_0 + \frac{\psi}{\eta} \cdot d_{jt} + \alpha^* k_j + \beta^* l_j + \gamma^*(m_j - \rho_t^s) + \delta^* G(n_j) + h(I_{jt}, k_{jt}, l_{jt}, z_{jt}) + \varepsilon_{jt}^*. \tag{15}$$

Here again $G(n)$ is assumed to be have the parametric form given by equation (12). We use this equation to estimate the import coefficient δ^* and the per-product import gain $a = \delta^*/\gamma^*$.

The following roadmap summarizes our empirical strategy of estimating these three equations.

1. Estimate $G(n)$ and S from firms' import shares (13).

2. Estimate γ^* from the material share (14).

3. Estimate a from (15) using the first step of the Olley-Pakes procedure.

4. Estimate α^*, β^*, η and ψ from (15) using the second step of the Olley-Pakes procedure.

We now turn to describe each step in more detail.

Estimating $G(n)$. Equation (13) links firms' import shares to the parametric $G(n)$ function. Figure 1 plots the average import share of firms as a function of n. Because few firms import many products we cannot precisely estimate the value $\bar{n}$ at which the curve flattens out, but the figure suggests that setting $\bar{n} = 150$ is a reasonable choice.[13] We take $\bar{G}$ from the input-output table as the share of non-service inputs among all intermediate inputs. We then estimate the curvature parameter λ and the import share S from (13) using a nonlinear least squares regression. To ensure that total imports computed in our estimated model using the estimated import shares aggregate up to total imports in the data, in this regression we weight observations by real intermediate spending. The graph also shows that the shape of our estimated $G(n)$ closely tracks the relationship between import share and the number of products.

Estimating the material coefficient γ^*. We estimate γ^* from the sample analogue of equation (14) as

$$\hat{\gamma}^* = \frac{\sum_j \sum_t M_{jt}}{\sum_j \sum_t R_{jt}}$$

which is the overall material share in the entire manufacturing sector.

Estimating the import effect using the first step of the Olley-Pakes approach. Equation (15) includes the demand proxy d_{jt}. This variable governs—conditional on industry demand q_t^s—the distribution of market shares within the industry. Exploiting spatial variation in demand, we proxy the demand shifter with local demand growth in the *county* of the firm. We compute this measure as the log growth in output of all firms (excluding j) in the year, industry, and county of firm j.

Equation (15) also includes the unknown control function $h(I_{jt}, k_{jt}, l_{jt}, z_{jt})$. We approximate this function with the sum of (i) a third-order polynomial of I_{jt}, k_{jt}, l_{jt} and the state variables o_{jt}, with coefficients that are allowed to differ by year; plus (ii) a linear function of industry by year effects, county effects, d_{jt} and q_t^s. These controls absorb several terms on

[13] Setting other plausible values for $\bar{n}$ had no effect on our qualitative results.

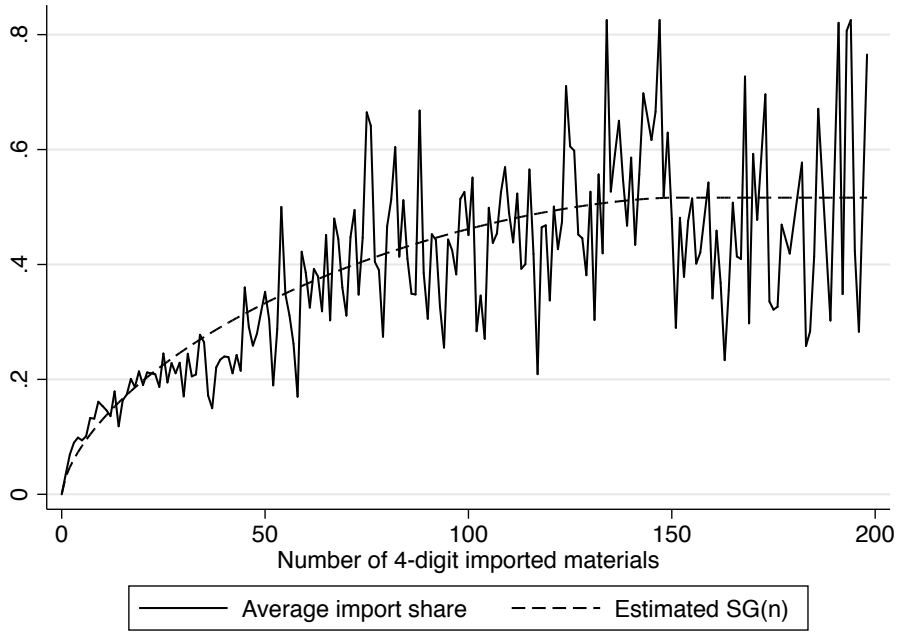

Figure 1: Import share as a function of the number of imported products

the right hand side of equation (15). Using our estimate of γ^* to subtract $\hat{\gamma}^*(m_{jt} - \rho_t^s)$ from both sides we obtain

$$r_{jt} - p_t^s - \gamma^*(m_{jt} - \rho_t^s) = \tilde{h}(I_{jt}, k_{jt}, l_{jt}, z_{jt}) + \delta^* G(n_j) + \varepsilon_{jt}^*. \tag{16}$$

where $\tilde{h}(I_{jt}, k_{jt}, l_{jt}, z_{jt})$ denotes the nonparameteric control function which absorbs ω_{jt}^{obs} as well as the capital, labor and demand effects. Because at this stage we already have an estimate of $G(n)$, and because ε is orthogonal to all remaining terms on the right hand side, we can estimate this equation with ordinary least squares to infer δ^*. We then use the formula $a = \delta^*/\gamma^*$ to estimate a.

Estimating the coefficients of state variables. Here we follow the second step of the Olley-Pakes procedure with the following modifications. (1) In our specification both k_{jt} and l_{jt} are predetermined, and productivity is also affected by the potentially persistent variables in o_{jt}. We estimate the coefficients of all of these variables in the second step. (2) Our revenue production function involves both q_t^s and d_{jt}, both of which are plausibly persistent. We thus treat them as state variables and estimate their coefficients in the second step. (3) Because we allow ω^{obs} to depend on industry and county effects, we also include these fixed effects in the second step.

Standard errors. We obtain standard errors for all estimates from a bootstrap.

We explain the precise implementation of these steps in more detail in Appendix B.

4.3 Logic of identification

The central identification problem associated with (9) is that ω is potentially correlated with the other determinants of output, including $G(n)$. We solve this problem using the Olley and Pakes approach, making structural assumptions that allow us to substitute out productivity. The identification of the import demand equation also follows from structural assumptions, which restrict the functional form on the right-hand side of (13). 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, a mechanism which could introduce spurious correlation between the share of imported inputs and $G(n)$ in (13). Importantly, our estimation is immune to this concern: productivity, which is explicitly incorporated in the model, cancels out of (13) 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 as a whole. 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 (13) 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. Because we start with (13), the form of $G(n)$ is determined as the shape traced out by the import share when n varies. Given this shape, S and A are estimated from the coefficients of $G(n)$ in the import share equation and in the revenue production function. Thus S and A are identified from variation in n given controls. In effect, we compare the output of two equally productive firms who import a different number of varieties. In the model, such variation in n comes from variation in the fixed costs, which affect the optimal number of imported inputs.

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 $G(n)$ in (13). We partially address this concern by explicitly allowing A—and hence a and S—to vary across groups of firms.

4.4 Recovering the quality and substitution parameters

Given our estimates of a and S we use the following two equations implied by the model to infer the deep parameters θ and A:

$$\theta = 1 - \frac{\log(1-S)}{a}, \tag{17}$$

$$\log A = a\left[1 - \frac{\log\ S}{\log(1-\ S)}\right]. \tag{18}$$

The basic idea behind these equations is that that a high per-product import gain a combined with a low import share S shows the importance of imperfect substitution. By (17), for a given gain from importing a, a lower S implies a lower θ: since importers are unwilling to switch to attractive foreign goods, the substitution elasticity must be low. And by (18), the gap between the quality effect $\log A$ and the total gain from imports a is a reflection of imperfect substitution, which is related to the import share S by (17).[14]

4.5 Extension to multiple groups

When different groups of firms, indexed by $g = 1, ..., \bar{g}$, have different efficiency of import use A_g, the import share S_g and the the per-product import gain a_g become group specific. This change affects the import share equation (13) which involves S_g and the revenue production function (15) which involves $\delta_g^* = \gamma^* a_g$. Our estimation strategy follows the same steps as above, but we now jointly estimate the group-specific parameters. Thus in estimating the import share equation we continue to use non-linear least squares, but in a specification which allows S_g to be group specific and requires that the shape of $G(n)$ is the same for all firms. Similarly, when estimating the revenue production function (15) we continue to use ordinary least squares, but allow δ_g to be group specific. We use this procedure to obtain preliminary estimates of a_g and S_g.

We next refine these estimates using the restrictions of our structural model to obtain our final, model-consistent estimates. To see why we do this, note that multiple groups also affect how we infer the quality and substitution parameters. Because A_g and S_g are group-specific, (17) and (18) become a system of equations. But because θ is the same for all

[14] Our approach here builds on 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, $x = 1 - S$ and the productivity gain is $\exp(a) = (1-S)^{1/(1-\theta)}$. Blaum, Lelarge and Peters (2014) derive a similar formula for the gains from input trade in a generalized version of our model in which firms can import from multiple countries.

firms, this system is over-identified. We therefore estimate A_g and θ using our preliminary estimates of a_g and S_g in a minimum distance procedure in which we minimize the sum of squared deviations across these equations. Then, given the deep parameters A_g and θ we invert equations (17) and (18) to obtain our refined, internally consistent estimates of the reduced-form parameters a_g and S_g. These are the estimates we report in the tables below.

4.6 Fixed costs

To estimate fixed costs, we make the additional assumption that the firm's fixed cost schedule depends only on its ownership status and on an i.i.d. random disturbance. Specifically, let the fixed cost of importing the n-th product by firm j in year t be $f_{jt}^n = \exp(\kappa_n^D 1_{jt}^D + \kappa_n^F 1_{jt}^F + \upsilon_{jt})$. Here 1_{jt}^F is an indicator for the firm having been foreign owned, $1_{jt}^D = 1 - 1_{jt}^F$; κ_n^F and κ_n^D measure the dependence of the fixed cost schedule on n, separately for firms that have or have not been foreign owned; and υ_{jt} is a shock that affects the entire fixed cost schedule. We assume that υ_{jt} is mean-zero, normally distributed, and independent of all other shocks, with a variance that may depend on the whether the firm has been foreign owned.

Recall from equation (11) that the optimal choice of n_{jt} is determined by trading off expected operating profits $\pi(n)$ with the fixed costs. Here $\pi(n)$ is a constant fraction of expected revenue, and by equation (10) the random variable over which expectations must be taken to compute it is $\exp(\varepsilon^*)$. We estimate this expected value as a weighted average of $\exp(\hat{\varepsilon}_{jt}^*)$ within industry s and year t, where $\hat{\varepsilon}_{jt}^*$ is our estimate of the realized error term. We choose weights to ensure that average expected revenue equals average actual revenue in each industry and year. Combining this measure with our coefficient estimates and the revenue production function (10) yields an empirical estimate of $\pi_{jt}(n)$ for each possible value of n.

From (11) the optimal choice of n is characterized by the inequalities $f_{jt}^n \leq \pi(n) - \pi(n-1)$ (for $n \geq 1$) and $f_{jt}^{n+1} > \pi(n+1) - \pi(n)$. Because we observe the actual choice of n_{jt} for all firms, given that the assumed distribution of υ_{jt} in our fixed cost model is normal, we can estimate this set of inequalities—separately for firms that have and have not been foreign owned—as an ordered probit. This allows us to infer the coefficients κ_n^F, κ_n^D, and the standard deviations σ_υ^F and σ_υ^D. These results yield an estimate for the distribution of the fixed cost vector $\bar{f}$. While we cannot infer its precise value, from the observed n_{jt} we can derive an interval bound on the realization of υ_{jt}. We then construct estimates for the fixed cost schedule of each firm in each year by calculating the conditional mean of the estimated distribution of υ within the interval bounds. This ensures that the implied choice of n_{jt} agrees with the choice we observe in the data. We use these estimates to compute summary

statistics of the fixed costs, and also in the counterfactual analysis. We explain the precise implementation of the fixed cost estimation in more detail in Appendix B.

5 Results

5.1 Basic results

Table 3 summarizes our basic results. For the production function parameters we report the "starred" values corresponding to the revenue production function, for example, $\alpha^* = \alpha(\eta - 1)/\eta$. From these, the quantity production function parameters can be recovered using our estimate of the demand elasticity η. 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.

Column 1 reports the results from our empirical procedure in a baseline specification in which the vector of additional state variables o_{jt} only includes an indicator for whether the firm has been foreign owned, and all firms have the same import quality parameter A. We estimate a highly significant per product import gain a of 0.33. This point estimate implies that the composite of the foreign and the domestic good is about $\exp(.33) - 1 = 39\%$ more efficient per dollar spent than the domestic good in itself. The share of non-service inputs among all intermediate inputs from the input-output table is $\bar{G} = 0.83$, and in column 1 the elasticity of output to intermediate inputs is estimated to be $\gamma^* = 0.75$. Combining these numbers, we predict that if a non-importer starts importing all tradeable varieties, it will experience an increase in log revenue productivity of $a\bar{G}\gamma^* = 0.20$, which corresponds to an increase in revenue productivity of about 22 percent.

The table also reports our estimates of the structural parameters A and θ. In the baseline specification, the price-adjusted quality advantage of foreign products relative to their domestic counterparts is $A = 1.19$. Based on our bootstrap (reported at the bottom of the table) this value value is different from 1 with $p = 0.004$. Imported inputs are thus about 19 percent better than domestic ones per dollar of expenditure. This difference in price-adjusted quality accounts for about 48 percent of the per-product import gain. The remaining 52 percent comes from imperfect substitution: we find that the elasticity of substitution between domestic and foreign goods is $\theta = 4$. 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.

Table 3: Baseline estimates

Dependent variable: log sales	Baseline (1)	Conditioning on exporter (2)	Conditioning on past imports (3)	Conditioning on lagged G(n) (4)
Capital (α*)	0.041 (0.003)	0.041 (0.003)	0.040 (0.003)	0.038 (0.004)
Labor (β*)	0.198 (0.008)	0.197 (0.008)	0.201 (0.008)	0.215 (0.015)
Materials (γ*)	0.752 (0.012)	0.752 (0.012)	0.752 (0.012)	0.752 (0.012)
Per-product import gain (a)	0.327 (0.063)	0.263 (0.058)	0.194 (0.053)	0.164 (0.055)
Import share (S)	0.626 (0.055)	0.626 (0.055)	0.626 (0.055)	0.626 (0.055)
Efficiency of imports (A)	1.186 (0.076)	1.147 (0.061)	1.107 (0.047)	1.089 (0.042)
Elasticity of substitution (θ)	4.006 [3.05; 6.07]	4.742 [3.52; 8.00]	6.053 [4.14; 12.92]	7.002 [4.45; 26.13]
Curvature of G(n) (λ)	0.650 (0.055)	0.650 (0.055)	0.650 (0.055)	0.650 (0.055)
Foreign owned	0.054 (0.013)	0.067 (0.010)	0.067 (0.010)	0.073 (0.126)
Exporter		0.046 (0.005)	0.046 (0.005)	0.052 (0.006)
Previous importer			-0.002 (0.006)	
Lagged per-product import gain				0.120 (0.045)
Industry sales (log)	0.066 (0.013)	0.066 (0.013)	0.031 (0.012)	0.076 (0.020)
Local demand growth	0.009 (0.006)	0.010 (0.006)	0.012 (0.005)	0.013 (0.007)
p-value of test for a=0	0.004	0.004	0.004	0.016
p-value of test for A=1	0.004	0.004	0.004	0.004
Observations	127,472	127,472	127,472	127,472

Notes: All specifications use the structural estimation procedure of Section 4. Bootstrapped standard errors clustered by firm are in parenthesis. For the elasticity of substitution (theta) we report a 95 percent confidence interval computed the same way in brackets.

The table also reports estimates of the demand terms. The coefficient of industry log sales (q_t^s) is 0.066, the inverse of which gives our estimate of the consumer's elasticity of substitution $\eta = 15$. Our estimate is towards the high end of the range of elasticities

reported in Broda and Weinstein (2006), and higher than De Loecker's estimates which range between 3 and 7. In particular, our estimate implies that the markup firms charge is about 6.7 percent of the price of the final good. We can use η to compute the parameters of the *quantity* production function. For example, our estimates imply that if a non-importer starts importing all tradeable varieties, it will experience an increase in log quantity productivity of $a\bar{G}\gamma \cdot \eta/(\eta - 1) = 0.22$, i.e., an increase in quantity productivity of 25 percent. The table also shows that the coefficient of local demand growth–our demand shifter—is a positive and significant 0.009. As expected, an increase in local demand maps into higher firm sales.

Column 2 re-estimates the model by also adding to the state variable vector o_{jt} an indicator for export market participation. 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 import estimates are somewhat smaller but similar to the previous specification ($a = 0.26$, $A = 1.15$), suggesting that our procedure succeeds in isolating the impact of imports on productivity.

Columns 3 and 4 explore the possibility that entering import markets entails a sunk, rather than a fixed cost. One variant of this hypothesis is that starting to import requires a sunk cost, but then importing additional varieties requires a per-period cost. To capture this force, in column 3 we include in o_{jt} an indicator for the past importing status of the firm. The estimated coefficients ($a = 0.19$ and $A = 1.11$) are smaller but still show a significant import effect.

In column 4 we explore another variant of the sunk cost hypothesis: that importing each additional variety requires a sunk cost. In this case the set of previously imported inputs becomes a state variable. To capture this mechanism, we include $\gamma G(n_{j,t-1})$ in o_{jt}. The estimated coefficients ($a = 0.16$ and $A = 1.09$) continue to show a significant import effect. Moreover, importing a good now also changes the set of available foreign goods and hence imports have *dynamic effects* on productivity. A simple way to measure these effects is with the coefficient of $\gamma G(n_{j,t-1})$.[15] This coefficient estimate shows that the "lagged per-product import gain" is 0.12. Combining this number with the (unlagged) per product import gain a, importing a product for two years would increase productivity by $\exp(0.12 + 0.164) - 1 = 33$ percent times the product's cost share.

The basic finding that imports have a positive productivity effect confirms and reinforces existing evidence from Indonesia (Amiti and Konings 2007), Chile (Kasahara and Rodrigue 2008) and India (Topalova and Khandelwal 2011). Our substantial per-product import gain

[15]We look at the coefficient of $\gamma G(n_{j,t-1})$ rather than that of $G(n_{j,t-1})$ to ensure that the coefficient is measured in the same units as a, which is the coefficient of $\gamma G(n_{j,t})$.

estimates support the analysis of Gopinath and Neiman (2013) who argue that the large drop in the number of inputs firms imported during a crisis in Argentina may have contributed to the observed decline in aggregate productivity.

A robust finding in columns 1-4 of the Table is that imperfect substitution is responsible for about half (between 48 and 52 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. And 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 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 evidence for this sort of interdependence in the context of combining foreign and domestic intermediate inputs.

5.2 Foreign ownership and the efficiency of import use

Firms that had been foreign owned played a very important role in the Hungarian economy. In our data the sales share of such firms in the manufacturing sector increased from 21 to 80 percent during 1992-2003. Moreover, across specifications in Table 3, firms that have been foreign owned are on average about 7 percent more productive than purely domestic firms, suggesting that growing foreign participation has had significant aggregate productivity effects in Hungary. A possibility is that foreign firms are more productive in part because the they use imports more efficiently. Indeed, these firms may have better access to low-cost input suppliers abroad, may have more extensive know-how about foreign goods, and may face lower transactions costs.

To explore this possibility, we implement our estimation procedure with two groups ($\bar{g} = 2$), allowing firms that have and have not been foreign owned to have different efficiency of import use A. Maintaining the assumption that firms use the same technology, the elasticity of substitution θ is held constant across all firms. Table 4 reports the results. The

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

Dep. variable: log sales	Baseline (1)		Conditioning on exporter (2)		Conditioning on past imports (3)		Conditioning on lagged G(n) (4)	
	Domestic	Foreign	Domestic	Foreign	Domestic	Foreign	Domestic	Foreign
Capital (α*)	0.041 (0.003)		0.041 (0.003)		0.040 (0.003)		0.039 (0.005)	
Labor (β*)	0.199 (0.008)		0.196 (0.008)		0.201 (0.008)		0.214 (0.019)	
Materials (γ*)	0.752 (0.012)		0.752 (0.012)		0.752 (0.012)		0.752 (0.012)	
Per-product import gain (a)	0.271 (0.063)	0.390 (0.068)	0.213 (0.057)	0.314 (0.063)	0.156 (0.052)	0.224 (0.059)	0.126 (0.062)	0.159 (0.064)
Import share (S)	0.490 (0.052)	0.621 (0.042)	0.486 (0.052)	0.625 (0.045)	0.490 (0.049)	0.621 (0.048)	0.515 (0.059)	0.596 (0.053)
Efficiency of imports (A)	0.984 (0.083)	1.220 (0.072)	0.982 (0.066)	1.178 (0.061)	0.991 (0.046)	1.121 (0.048)	1.010 (0.048)	1.071 (0.042)
Elasticity of substitution (θ)	3.484 [2.753;5.004]		4.118 [3.135;6.712]		5.323 [3.758;10.919]		6.719 [4.083;66.082]	
Curvature of G(n) (λ)	0.650 (0.055)		0.650 (0.055)		0.650 (0.055)		0.650 (0.054)	
Lagged per-product import gain (a)							-0.087 (0.055)	0.306 (0.095)
Foreign owned	0.066 (0.015)		0.061 (0.014)		0.066 (0.012)		0.028 (0.077)	
Exporter			0.045 (0.006)		0.046 (0.005)		0.052 (0.006)	
Previous importer					-0.002 (0.006)			
Industry sales	0.066 (0.013)		0.066 (0.013)		0.032 (0.012)		0.070 (0.020)	
Local demand growth	0.010 (0.006)		0.010 (0.006)		0.012 (0.005)		0.012 (0.007)	
p-value of test for a=0	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.219
p-value of test for A=1	0.777	0.004	0.733	0.004	0.777	0.004	0.897	0.046
p-value of test for a1=a2		0.012		0.016		0.012		0.269
Observations	127,472		127,472		127,472		127,472	

Notes: All specifications use the structural estimation procedure of Section 4. Bootstrapped standard errors clustered by firm are in parenthesis. For the elasticity of substitution (theta) we report a 95 percent confidence interval computed the same way in brackets.

regression results in the first specification show a large and significant difference in the per-product import gain. We obtain $a = 0.39$ for foreign and $a = 0.27$ for domestic companies, which come from differential ability of import use: we estimate $A = 1.22$ for foreign and $A = 0.98$, not significantly different from one, for domestic firms. These results imply that domestic companies benefit from imports primarily through imperfect substitution.[16]

The second specification adds the firm's exporter status to the vector of state variables o_{jt}. The qualitative results are as before: domestic firms are not better at using imported inputs than at using domestic inputs ($A_D = 0.98$), while foreign firms are ($A_F = 1.18$). The third specification also adds an indicator for past importing to the vector of state variables. Results are essentially unchanged. In the last specification, we also allow the lag of $G(n)$ to be a state variable. The per product import gain estimates are $a_D = 0.13$ and $a_F = 0.16$. This difference is smaller than before. However, the *lagged* per product import gain is much larger for foreign firms. Hence the long-run import gain—the sum of the current and the lagged per product gain—is significantly larger for foreign (0.41) than for domestic firms (0.11). We conclude that the efficiency advantage in using imports of firms that have been foreign owned is robust across specifications.

From a policy perspective it is important to understand whether their greater efficiency in import use A is *caused* by these firms having been 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 ("switcher firms"). In our sample there are 656 firms which switch from being domestically owned to being foreign owned. For each of these firms, using the baseline estimates from column 1 of Table 3, we compute residual quantity productivity as $\hat{\omega}_{jt} = \hat{\omega}^*_{jt}\hat{\eta}/(\hat{\eta} - 1)$, and residual revenue productivity as $\hat{\omega}^*_{jt} + \frac{1}{\hat{\eta}}q_{st} + \frac{\hat{\psi}}{\hat{\eta}}d_{jt}$. We calculate these measures as a function of event year, normalizing the date on which the firm becomes foreign-owned to zero, separately for firms that do import versus firms that do not import in that event year.

Figure 2 plots the difference between importers and and non-importers by event year, normalizing the difference to zero in the year before the acquisition, both for quantity and for revenue productivity. The Figure shows that after acquisition, both the quantity and the revenue productivity gap between importers and non-importers widens. The gap peaks

[16] As discussed in Section 4.5 the model-consistent estimates of a_g and S_g are filtered through the structural model of imperfect substitution and hence not equal to their preliminary (direct) estimates. But they are close: for example, $a_F = 0.39$ versus $a'_F = 0.38$ and $a_D = 0.27$ versus $a'_D = 0.29$ in the model-consistent respectively in the preliminary estimates.

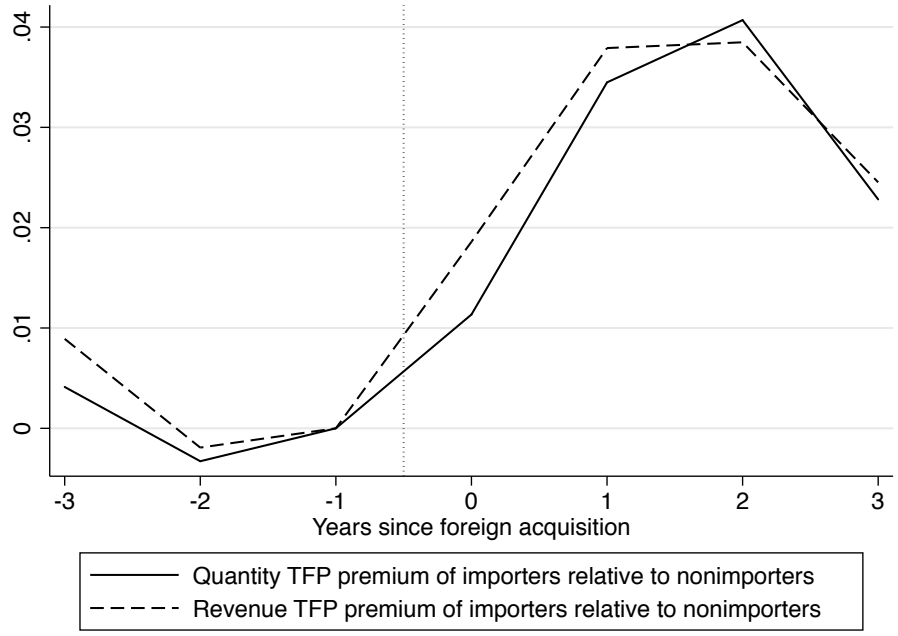

Figure 2: Productivity premium of importers around foreign acquisition

at about 4 percentage points for both measures in year 2 after the acquisition. The p-value that the gap widens is 0.11 for quantity and 0.06 for revenue productivity. These results are not fully conclusive, but they do suggest that part of the foreign premium in the efficiency of import use is causal. This in turn suggests a potential policy complementarity between financial and trade liberalization which we explore in Section 6 below.

5.3 Import effects by year and industry

To explore the robustness of our estimates and learn more about the impact of foreign goods, we next explore variation in the import effect over time. We estimate our model allowing for eight groups of firms ($\bar{g} = 8$) with potentially different efficiency of import use A. The groups are defined by whether the firm has been foreign owned, and by the three-year periods 1992-94, 1995-97, 1998-2000, and 2001-03. Table 5 reports the estimated A_g values. Our estimates are slightly noisier than before. But, consistent with the earlier findings, throughout the sample period imports have had a significant effect on productivity, and firms that have been foreign owned have been better in using imports. We will use these estimates to decompose productivity growth during 1993-2002 in Hungary into import-related and other channels in Section 6 below.

Table 5: The gains from importing over time

	Time period			
	1992-94	1995-97	1998-00	2001-03
Domestic firms				
Per-product import gain (a)	0.301	0.287	0.244	0.314
	(0.092)	(0.069)	(0.057)	(0.069)
Import share (S)	0.513	0.497	0.443	0.529
	(0.072)	(0.058)	(0.060)	(0.065)
Efficiency of imports (A)	1.022	0.996	0.909	1.049
	(0.126)	(0.095)	(0.099)	(0.110)
Foreign firms				
Per-product import gain (a)	0.332	0.459	0.401	0.364
	(0.089)	(0.092)	(0.072)	(0.072)
Import share (S)	0.548	0.667	0.617	0.582
	(0.064)	(0.058)	(0.057)	(0.055)
Efficiency of imports (A)	1.084	1.337	1.221	1.148
	(0.118)	(0.130)	(0.102)	(0.097)
p-value of domestic A=1	0.319	0.695	0.978	0.646
p-value of domestic A=previous column		0.420	0.424	0.214
p-value of foreign A=1	0.186	0.004	0.004	0.020
p-value of foreign A=previous column		0.113	0.404	0.271
Observations	127,472			

Notes: Table reports estimates of a single regression. Different coefficients for the productivity gain from importing (a), import share (S) and efficiency of imports (A) are estimated for foreign and domestic firms in each 3-year period using the structural procedure of Section 4. Other parameters are assumed to remain constant. Bootstrapped standard errors clustered by firm are in parenthesis.

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	18,977	0.602***	0.325	0.327**	1.654
18	Apparel	5,135	0.428***	0.811	1.453***	4.895
20	Wood products	7,926	0.363***	0.739	1.325	4.696
22	Printing and publishing	15,053	0.779***	0.701	1.733**	2.549
25	Rubber and plastics	7,473	0.405***	0.695	1.324***	3.926
26	Non-metallic minerals	5,203	0.508***	0.430	0.776	2.106
28	Fabricated metal products	17,782	0.356***	0.541	1.078	3.187
29	Machinery	13,966	0.350***	0.456	0.903	2.739
33	Instruments	5,232	1.495***	0.659	2.497	1.719

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.

We also examine how the import effect varies by industry. Because different industries might face different production possibilities and different market structures, 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. Because sectors differ substantially in the set of products they use, we also set by sector the value of n at which $G(n)$ reaches its maximum to be the 99th percentile of the number of imported products among all importers in that sector.

Table 6 reports our estimates of the key model parameters by industry.[17] 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 these gains on average. These results also highlight how the deep parameters are determined by our coefficient estimates. For example, the fabricated metal and the machinery industries have similar per-product import gains of 0.36 respectively 0.35. But because in the fabricated metal industry the import share is larger (54% versus 46%), our model implies a higher elasticity of substitution in that industry (3.2 versus 2.7). Intuitively, given the total advantage of foreign goods, a higher import share must come from greater substitutability.

5.4 Fixed costs

Table 7 reports summary statistics for the estimated fixed costs. This table uses the first specification in Table 4 which distinguishes the efficiency of import use for firms that have and have not been foreign owned. The top panel reports the median estimated fixed cost of importing the first product (f_{jt}^1) in four groups of observations: domestic nonimporters, domestic importers, foreign nonimporters and foreign importers. Importers have a much lower fixed cost, implying that there is selection into importing. Consistent with the idea that they have better connections with international suppliers, firms that have been foreign owned have lower fixed costs.

The bottom panel of the table reports the median cost of importing the *next*—that is, the $n_{jt} + 1$st—product, separately for domestic and foreign importers. Unlike what we saw in the top panel, this cost is higher for foreign importers. The reason for the difference can be understood by looking at Figure 3 which plots the estimated fixed cost schedule (setting

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

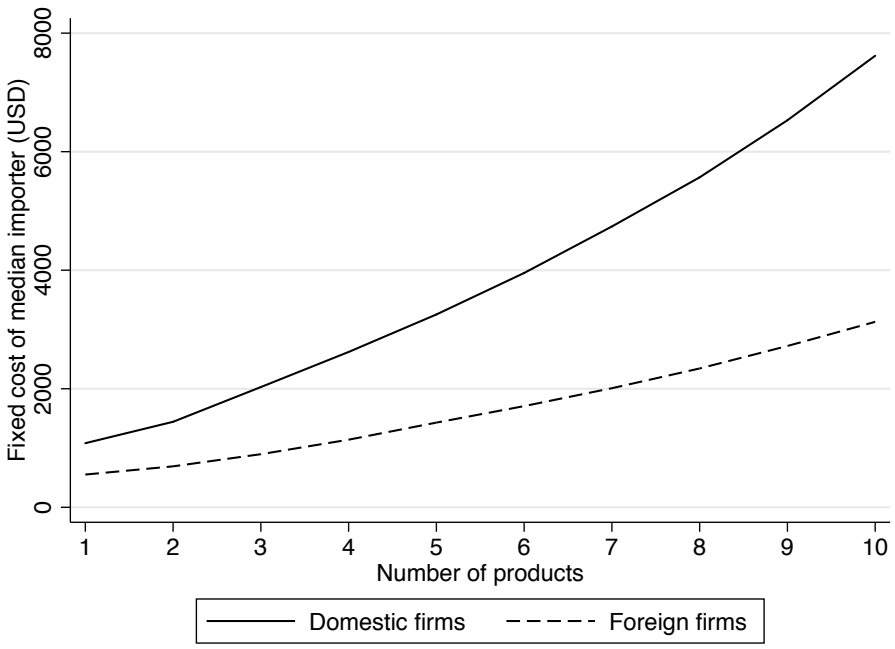

Figure 3: Fixed cost schedule of domestic and foreign firms

the firm-level disturbance $\upsilon_{jt} = 0$) as a function of n separately for firms that have and have not been foreign owned. Both schedules are increasing, but the one for foreign firms is below that for domestic firms. Consistent with the top panel in the table, for any given n foreign firms find it cheaper to import the n-th product. But because of the lower costs and the higher gain A they choose a higher n. Hence, consistent with the bottom panel in the table, firms that have been foreign owned face a higher fixed cost for the next product they could be importing.

It is helpful to understand the raw fact in the data that drives these findings. The fixed costs are estimated from the comovement between the model-implied gains from importing (which is related to firm size, A, and other factors) and the number of imported varieties n. The key fact is that n increases more steeply in the gain for foreign firms. Our structural model interprets this to show that foreign firms have lower fixed costs of importing.

Finally, we measure the extent to which accounting for fixed costs might affect our productivity estimates. 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 among all importers of their estimated fixed cost expenditures relative to their total production costs. We obtain 2.2%. In contrast, the average cost increase if firms produced the same

Table 7: Fixed costs

	Domestic	Foreign
Cost of importing first product ($)		
Nonimporter	35,938	14,890
Importer	1,097	556
Cost of importing next product ($)		
	4,165	6,429

Notes: Table reports the estimated median fixed costs of importing a good in 1998 US dollars, separately for domestic and foreign firms.

output without the use of imports, is, across all importers, 19% of production costs. Thus fixed costs amount to less than 12 percent of the cost savings created by imports. Ignoring them does not substantially alter the aggregate implications of our model.

6 Applications

This section develops two applications of our estimates. In Section 6.1 we quantify the aggregate productivity effects of imports in Hungary, and in Section 6.2 we explore the implications of tariff policies in our estimated economy.

6.1 Decomposing the productivity gains in Hungary

We decompose the growth in manufacturing productivity into various channels. To separate out the contribution of foreign ownership, we write a firm's residual log revenue productivity as $\omega^*_{jt} = 1^F_{jt}\mu^F_\omega + \chi_{jt}$. Here 1^F_{jt} is an indicator for whether the firm has been foreign owned, μ^F_ω is the Hicks-neutral mean log revenue productivity premium of foreign firms, and χ_{jt} measures remaining variation in revenue productivity, including demand effects. Also taking into account the effect of imports, the (log) revenue productivity of firm j in year t is

$$\phi^R_{jt} = [1^F_{jt} \cdot a_{Ft} + 1^D_{jt} \cdot a_{Dt}] \cdot \gamma^* G(n_{jt}) + 1^F_{jt}\mu^F_\omega + \chi_{jt}$$

where a_{Ft} and a_{Dt} denote the per product import gain in year t for firms that have and have not been foreign owned, and $1^D_{jt} = 1 - 1^F_{jt}$. Following Olley and Pakes (1996), we measure aggregate TFP as the sales-weighted average of firms' log TFP

$$\Phi^R_t = \sum_i \sigma_{jt}\phi^R_{jt}, \tag{19}$$

where σ_{jt} is the output share of firm j in year t. Denoting by $\tilde{G}^D_t$ and $\tilde{G}^F_t$ the sales-weighted average of $G(n_{jt})$ and by σ^D_t and σ^F_t the sales share of domestic and foreign firms in year t,

simple algebra shows that the growth in aggregate productivity between time t and time 0 equals

$$\begin{aligned}\Phi_t^R - \Phi_0^R = &[(a_{Dt} - a_{D0})\gamma^* \sigma_0^D \tilde{G}_0^D + (a_{Ft} - a_{F0})\gamma^* \sigma_0^F \tilde{G}_0^F] \\ &+ [a_{Dt}\gamma^* \sigma_0^D (\tilde{G}_t^D - \tilde{G}_0^D) + a_{Ft}\gamma^* \sigma_0^F (\tilde{G}_t^F - \tilde{G}_0^F)] \\ &+ [a_{Dt}\gamma^* (\sigma_t^D - \sigma_0^D)\tilde{G}_t^D + a_{Ft}\gamma^* (\sigma_t^F - \sigma_0^F)\tilde{G}_t^F] \\ &+ \mu_\omega^F (\sigma_t^F - \sigma_0^F) + \left[\sum_j (\sigma_{jt}\chi_{jt} - \sigma_{j0}\chi_{j0})\right].\end{aligned}$$

The five terms on the right-hand side have the following intuitive interpretation. The first two terms measure firm-level gains from 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 aggregate-level gains due to the increase in the share of foreign firms, which are more effective at using imports. Taken together, these terms represent the productivity gains created by importing. The fourth term measures the direct effect of increased foreign presence, caused by μ_ω^F, residual productivity premium of firms that have been foreign owned. And the final term measures additional, non-import-related variation in productivity. Note that in this decomposition firms that become foreign are assigned a higher efficiency of import use and a higher residual productivity. Thus the decomposition assumes that these gains are caused by foreign ownership.

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 the specification reported in Table 5. Although the specification in Table 5 is estimated in our main sample, 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 revenue production function (10) 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 the total growth in revenue productivity in the manufacturing sector in Hungary during this period was 21.1 percent. More than a quarter of this growth, 5.9 percent, can be attributed to various import-related mechanisms. Import-related gains at the firm level generated a productivity

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

Growth in aggregate productivity		21.1%
Coming from	intensive margin of imports	0.7%
	extensive margin of imports	4.0%
	increased import by foreign firms	1.2%
	direct effect of foreign firms	2.2%
	other	13.0%

Notes: Total productivity growth between the periods 1992-94 and 2001-03, interpreted as growth from 1993 to 2002, decomposed into the contributions of five different margins.

gain of 4.7 percent, most of which, 4 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 1.2 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.[18]

6.2 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. A reduction in tariffs directly increases productivity by reducing the price of imports and thereby increasing price-adjusted quality A. It also indirectly increases productivity 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 specification of Table 4. In particular, firms that have been foreign owned are more effective in using imports and also have higher residual productivity. The population of firms in the model economy is the union of all firm-year observations in our main sample.

[18]The remaining 15.2 percent of productivity growth is due to factors unrelated to importing. Higher foreign presence had a substantial direct effect by virtue of foreign firms being more productive; and we also find a 13 percent increase caused by forces outside our analysis.

To each firm we assign its actual capital, labor, foreign status, demand shifter and estimated ω^{obs} and ε. To each firm we also assign a fixed cost schedule realization, which we take to be a random draw from the estimated distribution of fixed costs, conditional on the interval bounds that ensure that the firm would optimally choose to import the observed number of imported varieties.[19] All firms face the same input prices, which we set exogenously to match aggregate output in each sector in each year. We then let firms make their optimal decisions about materials, imports, and the price of their final good, respecting the timing assumptions we had made in Section 4.

In the counterfactual experiments we track aggregate outcomes as firms' optimal choices change in response to changes in the environment. Because input prices are exogenous our experiments ignore general equilibrium price effects; and because capital and labor are exogenous they ignore dynamic considerations. We make these assumptions to keep the analysis simple and transparent.[20] Endogenizing capital and labor would likely amplify the effects we document due to reallocation. We describe the precise implementation of the counterfactual experiments in more detail in Appendix C.

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. 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 for both domestic and foreign firms.

Table 9 computes the change in aggregate productivity that results from a 10 percentage point reduction in tariffs in several hypothetical scenarios. Panel A focuses on the combination of decreasing tariffs and liberalizing FDI. The three columns correspond to environments which differ in the share of foreign-owned firms. In the first column we assume no firms are foreign, the second column is our benchmark economy, and in the third column we assume all firms are foreign. In keeping with the construction of the model economy, when we change the foreign status of a firm, we change its efficiency of using imports A, adjust its log fixed cost schedule by the mean difference in log fixed cost schedules of foreign and domestic firms, and adjust its Hicks neutral log productivity by the foreign productivity premium.

The middle column in Panel A shows that the effect of a tariff cut is nonlinear. As the first row shows, a tariff reduction from 40 percent to 30 percent increases aggregate

[19] We use a random draw rather than the mean conditional on the interval bounds to make the distribution of fixed costs in the model economy smooth.

[20] The partial equilibrium approach can be justified if the policies we evaluate affect a small industry in a small open economy.

Table 9: Counterfactual experiments

Panel A

Tariff reduction	No firms foreign	Baseline	All firms foreign
40% to 30%	0.8%	1.3%	1.6%
10% to 0%	1.6%	2.5%	2.9%

Panel B

Tariff reduction	High fixed cost	Baseline	Low fixed cost
40% to 30%	1.2%	1.3%	1.5%
10% to 0%	2.2%	2.5%	2.7%

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.

productivity by 1.3 percent. In contrast, as the second row shows, a tariff reduction from 10 percent to 0 percent increases aggregate productivity by 2.5 percent. That tariff cuts have larger effects in a more open economy may 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. This logic 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.

Comparing across columns in Panel A also reveals a policy complementarity between FDI liberalization and trade liberalization. When no firms are foreign, reducing tariffs from 40 percent to 30 percent has a 0.8 percent productivity effect. When all firms are foreign, the same tariff cut has a 1.6 percent productivity effect. This complementarity emerges because foreign firms are more effective in using imports. As the second row shows, this complementarity is slightly stronger in a more open economy.

Panel B of Table 9 focuses on the combination of decreasing tariffs and changing the fixed costs—such as those associated with licensing—of importing.[21] In the high fixed cost column each firm is assigned three times its baseline fixed cost vector; in the middle column each firm is assigned its baseline fixed cost vector; and in the low fixed cost column each firm is assigned one-third of its baseline fixed cost vector.[22]

[21] In related work, Hornok and Koren (2015) explore the effect of changing fixed costs on trade flows.

[22] 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

We find that tariff effects are larger with lower fixed costs. In the first row, a tariff cut from 40 percent to 30 percent increases productivity by 1.2 percent in the high fixed cost environment, and by 1.5 percent in the low-fixed-cost environment. In a more open economy, these effects are larger. In the second row, a tariff cut from 10 percent to zero increases productivity by 2.2 percent in the high fixed cost environment, and by 2.7 percent in the low-fixed-cost environment. These results point at a policy complementarity between reducing tariffs and reducing the fixed costs of importing, and suggest that this complementarity is stronger in a more open economy.

Our results 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 fixed cost 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 an increase in firm productivity (Topalova and Khandelwal 2011). And consistent with the FDI complementarity, these effects were stronger in industries with higher FDI liberalization (Topalova and Khandelwal 2011).

Import substitution. Finally we 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. For simplicity we perform this analysis in a model in which foreign and domestic firms use imports equally efficiently, taking the parameters from column 1 of Table 3. 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) As in our baseline results, about 52 percent of the gains are due to imperfect substitution. We implement these scenarios by holding fixed the per-product import gain $a = 0.33$, and by adjusting A and θ for the different scenarios. For instance, in the first scenario we set $\theta = 20$ which implies effectively perfect substitution, and let $A = \exp(a)$.

Figure 4 plots, as a function of the tariff level, the log dollar value of domestic input use in these scenarios. Values are measured relative to the baseline model with zero tariffs. Begin with the curve corresponding to the first scenario, in which the import effect comes only from quality differences. In this case domestic import demand is initially flat and then rapidly increasing. In contrast, the curve corresponding to no quality differences has a uniform small slope. This difference is intuitive: when foreign goods are perfect substitutes, there exists

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.

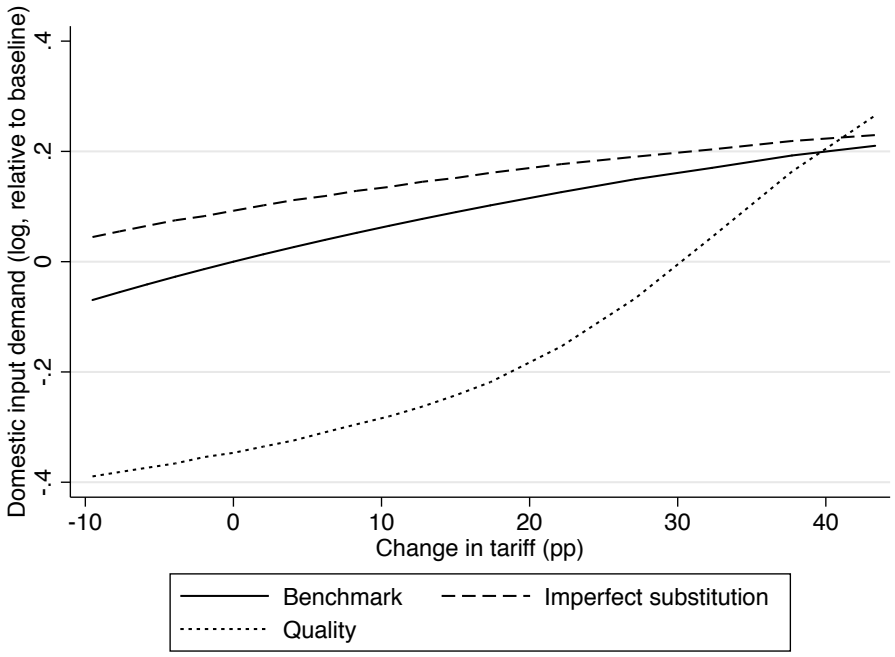

Figure 4: Effect of tariff changes on domestic input demand

a range in which small price changes 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 also relatively flat. This curve also reflects the effect that the losses caused by the tariff cut are counteracted by increased demand for all inputs created by the productivity gains from importing. A key lesson from the figure is that the magnitude of redistributive losses due to import substitution depend strongly on the extent of substitution and on the initial level of tariffs. More broadly, 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.

7 Conclusion

This paper explored the effect of imports on productivity by estimating a structural model of importers in a panel of Hungarian firms. We found that imports have a significant and large effect on firm productivity, about half of which is due to imperfect substitution between foreign and domestic goods. We also found that foreign firms use imports more effectively and pay lower fixed import costs. We then used our estimates in combination with our structural model to conduct counterfactual analysis. This analysis showed that during 1993-

2002, a third of the productivity growth in Hungary was due to imported inputs. It also showed that the productivity gain from a tariff cut is larger when the economy has many importers and many foreign firms, implying policy complementarities between tariff cuts, dismantling non-tariff barriers, and FDI liberalization.

Perhaps the main caveat to our analysis is that, in the absence of exogenous variation, we need to use with full force the restrictions imposed by our structural framework. However, a benefit of our structural framework is that it allows for explicit counterfactual analysis. Our framework and analysis may be extended in a number of ways. One possibility is to seek reduced-form evidence for our new predictions, such as those concerning policy complementarities. A second direction is to use our formal model to examine concrete episodes—such as crises, as explored by Gopinath and Neiman (2013)—in which the imported goods margin is relevant. A third direction is to extend our framework to also incorporate capital goods. Work by Caselli and Wilson (2004) suggests that, because of the technology embedded in them, capital imports can have a substantial effects on productivity. Investigating these directions can improve 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.

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.

Bernard, A. B., Jensen, J. B., Redding, S. J. and Schott, P. K. (2012). The empirics of firm heterogeneity and international trade, *Annual Review of Economics* **4**(1): 283–313.

Blaum, J., Lelarge, C. and Peters, M. (2014). Estimating the productivity gains from importing, Working paper, Brown University.

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.

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.

De Loecker, J. (2011). Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity, *Econometrica* **79**(5): 1407–1451.

De Loecker, J., Goldberg, P. K., Khandelwal, A. K. and Pavcnik, N. (2014). Prices, markups, and trade reform, Working paper, Princeton University, Yale University, Columbia University and Dartmouth College.

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.

Gandhi, A., Navarro, S. and Rivers, D. (2013). On the identification of production functions: How heterogeneous is productivity?, Working paper, University of Wisconsin-Madison and University of Western Ontario.

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.

Gopinath, G. and Neiman, B. (2013). Trade adjustment and productivity in large crises, *American Economic Review* **104**(3): 793–831.

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.

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

Hornok, C. and Koren, M. (2015). Administrative barriers to trade, *Journal of International Econonomics* **in press**.

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

A.1 Deriving the revenue production function

Demand. Constrained maximization of the utility function (3) implies

$$Q^{1/\eta} V_j^{1/\eta} Q_j^{-1/\eta} = \lambda P_j \tag{20}$$

where λ is the Lagrange multiplier of the budget constraint. Multiplying by Q_j and then summing over all firms in industry s, using the notation that firm revenue is $R_j = P_j Q_j$ and industry revenue is $R = \sum_{j=1}^{J_s} R_j$, and recalling that equation (3) also defines industry quantity Q, we obtain $Q = \lambda R$. Based on this we can define the industry price index as $P = 1/\lambda$, plug this back into the first order condition (20) and raise that to the power $-\eta$ to obtain the demand for the product of firm j

$$\frac{Q_j}{Q} = V_j \left(\frac{P_j}{P} \right)^{-\eta}. \tag{21}$$

Multiplying by P_j and summing over j now gives the familiar expression $P^{1-\eta} = \sum_j V_j P_j^{1-\eta}$ for the industry price index.

We then use (21) to express firm revenue, deflated by the industry price index, with firm and industry output as

$$\frac{R_j}{P} = Q^{1/\eta} V_j^{1/\eta} Q_j^{(\eta-1)/\eta}. \tag{22}$$

Production function. Given the Cobb-Douglas structure, total expenditure on intermediates M must equal their price index times their Cobb-Douglas aggregate:

$$M_j = \prod_{i=1}^{N} (\gamma_i/\gamma)^{-\gamma_i/\gamma} \prod_{i=1}^{N} P_{ji}^{\gamma_i/\gamma} \prod_{i=1}^{N} X_{ji}^{\gamma_i/\gamma}. \tag{23}$$

By (4) and (5), $P_{ji} = P_{iH}\exp(-a)$ for $i \leq n_j$ and $P_{ji} = P_{iH}$ otherwise. Denoting the first term in (23) by Γ, we have

$$M_j = \Gamma \prod_{i=1}^{N} P_{iH}^{\gamma_i/\gamma} \prod_{i=1}^{n_j} \exp(-a\gamma_i/\gamma) \prod_{i=1}^{N} X_{ji}^{\gamma_i/\gamma}.$$

It follows that

$$M_j = \exp[-aG(n_j)]\Gamma \prod_{i=1}^{N} P_{iH}^{\gamma_i/\gamma} \prod_{i=1}^{N} X_{ji}^{\gamma_i/\gamma}$$

and

$$\prod_{i=1}^{N} X_{ji}^{\gamma_i} = M_j^{\gamma} \exp[a\gamma G(n_j)]\Gamma \prod_{i=1}^{N} P_{iH}^{-\gamma_i}. \tag{24}$$

Define the industry input price index as

$$\varrho = \Gamma \prod_{i=1}^{N} P_{iH}^{\gamma_i/\gamma},$$

the (share-weighted) geometric average of domestic input prices. We assume that this is the input price index reported by the statistical office. The constant Γ only pins down the level of prices, and hence does not affect the price index.

Taking logs in (24), substituting in $\rho = \log(\varrho)$ and combining the result with (1) yields

$$q_j = \alpha k_j + \beta l_j + \gamma(m_j - \rho) + aG(n_j) + \omega_j$$

which is the quantity production function (9). Combining it with (22) yields

$$r_j - p = \frac{1}{\eta} q + \frac{1}{\eta} v_j + \alpha^* k_j + \beta^* l_j + \gamma^*(m_j - \rho) + \gamma^* aG(n_j) + \omega_j^*$$

which is the revenue production function (10).

A.2 Profits as a function of the number of imported inputs

We now compute operating profits as a function of the number of imported inputs n, assuming that other freely adjustable inputs are chosen optimally. Spending on intermediate inputs is chosen before ε_j is realized. Because of the Cobb-Douglas structure, intermediate spending is a constant γ^* share of expected revenue $M_j = \gamma^* E_\varepsilon(R_j)$, and expected operating profits are the remaining share $\pi_j(n) = (1-\gamma^*)E_\varepsilon(R_j)$. Because capital and labor had been chosen in advance, their costs are sunk at this stage. Substituting in $M_j = \gamma^* E_\varepsilon(R_j)$, the revenue production function (10) implies that

$$E_\varepsilon(R_j) = V_j^{1/\eta} P Q^{1/\eta} E\left(e^{\varepsilon_j^*}\right) K_j^{\alpha^*} L_j^{\beta^*} \left(\frac{\gamma^* E_\varepsilon R_j}{\varrho}\right)^{\gamma^*} e^{\gamma^* a G(n_j)} \Omega_j^*. \tag{25}$$

From this we can compute expected revenue for a firm that does not import by setting $n_j = 0$ and rearranging the equation to solve for $E_\varepsilon(R_j)$ which appears on both sides. Since variable profits are a fraction $1 - \gamma^*$ of expected revenue we then obtain

$$\pi_j(0) = (1-\gamma^*)\left(V_j^{1/\eta} P Q^{1/\eta} E\left(e^{\varepsilon_j^*}\right) K_j^{\alpha^*} L_j^{\beta^*} \left(\frac{\gamma^*}{\varrho}\right)^{\gamma^*} \Omega_j^*\right)^{1/(1-\gamma^*)}.$$

Combining (25) and $\pi_j(0)$ also gives expected operating profits from importing n varieties

$$\pi_j(n) = \pi_j(0) \exp\left(\frac{\gamma^* a}{1-\gamma^*} G(n)\right).$$

B Estimation

B.1 Estimating the coefficients

First step of the Olley-Pakes procedure. We implement the estimation of (16) by first regressing both sides on the flexible controls $\tilde{h}(I_{jt}, k_{jt}, l_{jt}, z_{jt})$, then taking the residuals, and then estimating the regression on the residuals using ordinary least squares. We follow this approach because it is computationally easier, and because the coefficients of the terms in $\tilde{h}(I_{jt}, k_{jt}, l_{jt}, z_{jt})$ are not of direct interest to us.

Second step of the Olley-Pakes procedure. Recall that we can write $\omega_{jt}^{obs} = \mu(s,c,o) + \varpi_{jt}$ where $\mu(s,c,o) = \mu_s^1 + \mu_c^2 + \mu_o \cdot o$ and ϖ_{jt} is a Markov process satisfying $\varpi_{jt} = f(\varpi_{j,t-1}) + e_{jt}$ where e_{jt} are i.i.d. and independent of all other shocks.

Building on Olley and Pakes (1996) we regress exit in t on (i) a third-order polynomial of $I_{j,t-1}$, $k_{j,t-1}$ and $l_{j,t-1}$ and the lagged variables $o_{j,t-1}$, with coefficients that are allowed to differ by year; plus (ii) a linear function of (lagged) industry by year effects and county effects, and $d_{j,t-1}$ and q_{t-1}^s. We denote the predicted exit probability by $\hat{p}_{jt}^{exit}$. Then $\hat{p}_{jt}^{exit}$ and $\varpi_{j,t-1}$ provide sufficient statistics about the bias in ϖ_{jt}: denoting information available at $t-1$ by Info_{t-1}, we have

$$E(\varpi_{jt} | \text{exit}_{jt} = 0, \text{Info}_{t-1}) = \psi(\hat{p}_{jt}^{exit}, \varpi_{j,t-1}). \tag{26}$$

Because we do not observe ϖ_{jt} or $\varpi_{j,t-1}$, we express them using both the already estimated and the as yet unknown production function coefficients. Denote by ζ the vector of (un-

known) coefficients α, β, η, μ_t^1, μ_c^2, μ_o. Then from (10) we get

$$\hat{\varpi}_{jt}(\zeta) = \left(r_{jt} - p_t^s - \gamma^*(m_{jt} - \rho_t^s) - \delta^* G(n_{jt}) - \varepsilon_{jt}^*\right) - \left(\frac{1}{\eta}q_t^s + \frac{\psi_0}{\eta} + \frac{\psi}{\eta}d_{jt} + \alpha^* k_{jt} + \beta^* l_{jt} + \mu_s^{1*} + \mu_c^{2*} + \mu_o^* \cdot o\right)$$

where the unknowns in the first parenthesis were estimated in the first step—in particular ε_{jt}^* is the residual—and the coefficients in the second parenthesis form the ζ vector which needs to be estimated. Note that ψ_0/η cannot be separately identified from the fixed effects, so we ignore it.

We then estimate ζ from (26) which, for firms that do not exit in year t, we can write as

$$\varpi_{jt}(\zeta) = \psi(\hat{p}_{jt}^{exit}, \varpi_{j,t-1}(\zeta)) + \tilde{e}_{jt}$$

where ψ is the conditional expectation and hence $\tilde{e}_{jt}$ is uncorrelated with past productivity and exit. To estimate this equation we approximate ψ as a third-order polynomial, and impose the moment condition that $\tilde{e}_{jt}$ is orthogonal to the following set of instruments: $\hat{p}_{it}^{exit}$, $\hat{\varpi}_{j,t-1}(\zeta)$, k_{jt}, l_{jt}, $k_{j,t-1}$, $l_{j,t-1}$, o_{jt}, q_t^s, $q_{s,t-1}$ and d_{jt}.

Implementing this estimation using GMM is computationally difficult because the many fixed effects μ_s and μ_c substantially increase the dimensionality of the maximization problem. We therefore use an approximation to infer these fixed effects and estimate only the remaining components of ζ with GMM. Specifically, we estimate μ_s and μ_c as the average of observed productivity ω_{jt}^{obs} across all firms and years for the given industry and county. Because we only observe surviving firms, this approach yields biased estimates of the unconditional means. However, if the bias is the same across industries and counties, then it would be subsumed in the regression constant. To check whether this is the case, we explore to what extent exit rates vary between industries and counties. Industry and county fixed effects explain only 0.6% of the variation in exit rates across firms. This suggests that the heterogeneity in the bias which we ignore by using simple means is likely to be small.

B.2 Bootstrap

We obtain standard errors, confidence intervals and p-values from a bootstrap with 500 draws. We sample firms with replacement, holding their entire time path together to preserve the joint distribution of variables at different points in time. This is akin to clustering standard errors by firm. We then estimate the model for each of the 500 draws. Standard errors are obtained as the empirical standard deviation of the 500 estimates. We define the 95 percent

confidence interval of a parameter (reported for example for the elasticity of substitution θ) as the range between the 2.5th and 97.5th percentile of the empirical distribution of the estimates for that parameter.

We conduct the two-tailed test for the hypothesis that $A = 1$ as follows. If the estimate $\hat{A} > 1$, we count the number of realizations in the bootstrap for which $\hat{A}_{(bs)} \leq 1$ and we compute the p-value as $2/500$ times this number. Similarly, if $\hat{A} < 1$, the p-value is based on the number of estimates for which $\hat{A}_{(bs)} \geq 1$.

B.3 Fixed costs

Estimating the parameters. Denoting the number of observations in industry s and year t by N_{st}, we estimate the expected unobserved productivity shock as a weighted average across firms,

$$\hat{E}_{st}\left(e^{\varepsilon^*}\right) = \frac{\sum_{j=1}^{J_{st}} \Phi_{jt} e^{\varepsilon^*_{jt}}}{\sum_{j=1}^{J_{st}} \Phi_{jt}}$$

where the right hand-side sums over all observations in industry s and year t. The weight

$$\Phi_{jt} = \exp\left[\frac{1}{\eta} q_t^s + \frac{1}{\eta}\psi_0 + \frac{\psi}{\eta} \cdot d_{jt} + \alpha^* k_j + \beta^* l_j + \gamma^*(m_j - \rho_t^s) + \delta^* G(n_j) + h(I_{jt}, k_{jt}, l_{jt}, z_{jt})\right]$$

is chosen based on (15) to ensure that total actual industry revenue equals total expected industry revenue in each s and t.

For firm j in year t, we denote by

$$\Delta\pi_{jt}(n) = \pi_{jt}(0)\left[e^{\frac{\gamma^* a}{1-\gamma^*}G(n)} - e^{\frac{\gamma^* a}{1-\gamma^*}G(n-1)}\right]$$

the increase in variable profit from importing the nth product. We introduce the notation

$$\nu_n = \ln\left[e^{\frac{\gamma^* a}{1-\gamma^*}G(n)} - e^{\frac{\gamma^* a}{1-\gamma^*}G(n-1)}\right]$$

so that

$$\Delta\pi_{jt}(n) = \pi_{jt}(0)e^{\nu_n}.$$

The firm chooses to import n_{jt} products if

$$f_{jt}^{(n_{jt})} \leq \Delta\pi_{jt}(n_{jt}), \text{ and } f_{jt}^{(n_{jt}+1)} > \Delta\pi_{jt}(n_{jt}+1).$$

Recall that we model the fixed cost schedule as $f_{jt}^{(n)} = \exp(\kappa_n^D 1_{jt}^D + \kappa_n^F 1_{jt}^F + \upsilon_{jt})$. We estimate the fixed costs separately for firms that have or have not been foreign owned. Consider firms that have been foreign owned ($1_{jt}^F = 1$, $1_{jt}^D = 0$). Taking logs of the above inequalities,

substituting in the fixed costs and ν_n, and rearranging, we obtain the following optimality conditions

$$\kappa^F_{n_{jt}} - \nu_{n_{jt}} \leq \ln \pi_{jt}(0) - \upsilon_{jt} < \kappa^F_{n_{jt}+1} - \nu_{n_{jt}+1}. \tag{27}$$

We use *ordered probit* to estimate the parameters of these inequalities. More specifically, dividing by σ^F_υ yields

$$\frac{1}{\sigma^F_\upsilon}(\kappa^F_{n_{jt}} - \nu_{n_{jt}}) \leq \frac{1}{\sigma^F_\upsilon} \ln \pi_{jt}(0) - \frac{1}{\sigma^F_\upsilon}\upsilon_{jt} < \frac{1}{\sigma^F_\upsilon}(\kappa^F_{n_{jt}+1} - \nu_{n_{jt}+1}).$$

This is exactly a set of ordered probit inequalities with $\frac{1}{\sigma^F_\upsilon} \ln \pi_{jt}(0) - \frac{1}{\sigma^F_\upsilon}\upsilon_{jt}$ as the latent variable and $\frac{1}{\sigma^F_\upsilon}(\kappa^F_{n_{jt}} - \nu_{n_{jt}})$ as bounds. The coefficient of $\ln \pi_{jt}(0)$ is $1/\sigma^F_\upsilon$, which gives us an estimate of σ^F_υ. Combining this parameter with the ν_n which we can express as a function of our estimates of a, γ^* and $G(n)$, we can recover the fixed cost function κ_n. A similar set of inequalities holds for firms that have not been foreign owned, and we estimate the parameters of those analogously.

Estimating the fixed cost realization. The total importing cost paid by a firm that has been foreign owned and imports n_{jt} products is

$$\sum_{i=1}^{n_{jt}} f^{(i)}_{jt} = e^{\upsilon_{jt}} \sum_{i=1}^{n_{jt}} e^{\kappa^F_i}.$$

We do not know the realization of υ_{jt}, but we can put bounds on it based on the profit maximization inequalities (27):

$$\ln \pi_{jt}(0) + \nu_{n_{jt}+1} - \kappa^F_{n_{jt}+1} < \upsilon_{jt} \leq \ln \pi_{jt}(0) + \nu_{n_{jt}} - \kappa^F_{n_{jt}}.$$

Here we have estimates for both the upper and the lower bound. We can then draw a simulated fixed cost by drawing an υ_{jt} from its estimated distribution conditional on these bounds. The $\bar{f}_{jt}$ constructed from that υ_{jt} is our estimate of the firm's fixed cost schedule. We use the same procedure to estimate the fixed cost schedule of firms that have not been foreign owned.

C Counterfactuals

We conduct counterfactual experiments with respect to (1) tariffs τ, (2) the fixed costs of importing f, and (3) the share of foreign firms in the economy. We now discuss how we measure the effect of these counterfactual changes on total factor productivity and the demand for domestic inputs.

Tariffs affect relative prices in equation (4). We assume that both the domestic and the pre-tax import prices are unaffected by the tariff change. Then, a tariff rate of $\tau > 0$ raises the relative price of imported goods by a factor $1+\tau$. The productivity gain from importing a particular product is now smaller,

$$a' = \frac{\log\left[1 + \left(\frac{A}{1+\tau}\right)^{\theta-1}\right]}{\theta - 1}. \tag{28}$$

Given these counteractual parameters, we use equation (11) to solve for the new optimal number of imported products n'. Because the per product import gain is reduced, firms will typically enter fewer import markets. The overall (revenue) productivity gains from importing are reduced to $\gamma^* a' G(n')$. We use these values to compute both firm-level and aggregate productivity in the counterfactual.

As both the number of imported products and the amount imported from each product fall, overall demand for imports also fall. There is a corresponding increase in the demand for domestic products. We use equation (8) and (14) to express the domestic input demand of firm j as

$$M_j^{H\prime} = \gamma^* R_j' S' G(n_j'),$$

and then add up across firms to get the total demand.

Varying the fixed cost schedule $\bar{f}$ has no direct effect on a or S, but it does affect the costs in the optimization problem (11) for the number of imported varieties. We account for these costs and otherwise proceed as above in this counterfactual.

Changing the ownership status of the firm to foreign involves (i) changing the per-product import gain from a_D to a_F; (ii) changing the fixed cost schedule to that estimated for foreign firms; and (iii) a direct change in mean productivity by μ_ω^F. After implementing these changes we proceed as above in computing the outcomes of interest.