

From Groundnuts to Globalization: A Structural Estimate of Trade and Growth*

June 9, 2006

Christian Broda, Joshua Greenfield, and David E. Weinstein

Abstract

Starting with Romer [1987] and Rivera-Batiz-Romer [1991] economists have been able to model how trade enhances growth through the creation and import of new varieties. In this framework, international trade increases economic output through two channels. First, trade raises productivity levels because producers gain access to new imported varieties. Second, increases in the number of varieties drives down the cost of innovation and results in ever more variety creation. Using highly disaggregate trade data, e.g. Gabon's imports of Gambian raw, unshelled, groundnuts, we structurally estimate the impact that new imports have had in approximately 4000 markets per country. We then move from groundnuts to globalization by building an exact TFP index that aggregates these micro gains to obtain an estimate of trade on productivity growth for each country. We find that in the typical country in the world, new imported varieties account for 12 percent of its productivity growth. These effects are larger in developing countries where the median impact new imported varieties equals a quarter of national productivity growth.

* Christian Broda is at the University of Chicago, Graduate School of Business; Joshua Greenfield is at Columbia University, and David Weinstein is at Columbia University and the NBER. Broda and Weinstein are grateful to the National Science Foundation for its support of this research (NSF grant #0214378). We wish to thank Andres Rodriguez-Clare, Robert Feenstra, and Elhanan Helpman for excellent comments.

I. Introduction

Economists have long postulated that trade may raise growth. However, it was not until the work of Romer [1987] and Rivera-Batiz and Romer [1991] that we had a general equilibrium model that would let us understand how trade might bring this about. While this seminal work has spawned the development of the vast endogenous growth literature, it has fallen short of one of its main objectives. As Rivera-Batiz and Romer state in their opening paragraph, they did not write their article because no one had thought that new traded varieties could foster growth – after all, the idea was already widely believed at the time – but rather because “it would be difficult for any of us to offer a rigorous model that has been (or even could be) calibrated to data.” Seen in this light, the failure of anyone to calibrate or structurally estimate their model (or even related models, e.g. Jones [1995]) means that while most economists continue to believe that globalization raises growth, we still know precious little about the magnitudes or mechanisms (c.f. Easterly and Levine [2001]).

A major reason for this failure stems from the difficulty of bridging the gap between micro and macro evidence on the effects of trade on growth. Trade and development economists have estimated rich models of what happens to particular firms and individuals in particular cases of liberalization, but it is very hard to generalize from these detailed econometric case studies. Empirical macroeconomic studies, by contrast, have tended to use “one regression fits all” specifications to examine how globalization affects growth around the world, but their grand assumptions leave us with little to say about the precise mechanisms underlying the results. It is as if the only way that economists can view the world is through microscopes and telescopes.

As a result, these two strands of literature try to describe the same processes, but they have little influence on each other, and a deep skepticism has developed regarding the robustness of the links between trade and growth. For example, in the excellent survey by Hallak and Levinsohn [2004], the authors identify three main classes of “basic methodological shortcomings” in the cross-country evidence. First, trade policy or openness is typically summarized by a one-dimensional index that has little theoretical foundation. Second, there are severe omitted variables biases, which lead to results that are not robust (c.f. Sala-i-Martin [1997], Rodriguez and Rodrik [2001], Nogueer and Siscart [2005, and especially 2006]). Finally,

there is so much heterogeneity in economic conditions across countries that it is doubtful that there is a unique mapping of trade into growth.

The problems in the cross-country growth regressions lead Hallak and Levinsohn [2004] to conclude that “more progress might be made by asking ‘smaller’ questions and writing papers such as ‘The impact of trade liberalization on groundnut farmers in Senegal and Gambia.’” These “micro-econometric studies” constitute a second approach to understanding globalization. The enormous advantage of these studies is that they provide rich and compelling econometric case studies of particular liberalizations. In the best examples, one can often be extremely precise about what exactly is being estimated and how the pieces fit together. The disadvantage is that it is very hard to extrapolate from groundnuts to globalization.

This paper can be thought of as a hybrid approach to understanding how trade affects growth. By breaking world trade down into 6-digit bilateral import flows and estimating hundreds of structural parameters per country, we are able to build estimates that preserve the cross-country and cross-industry richness of the global economy. The estimation strategy we use enables us to employ a theoretically justified aggregation structure to understand the gains from new varieties for each country. We can use these estimates to account for productivity growth predicted by endogenous growth models and thereby can be very precise about the channel through which new imported varieties affect growth.

The first part of the paper focuses on what Rivera-Batiz and Romer call the “level effect,” i.e. the productivity gain arising from the import of new and better varieties. These gains are predicted by static and dynamic models that have production functions based on the Dixit-Stiglitz [1977] framework.¹ One of the most important stylized facts about economic growth is that productivity growth, rather than factor accumulation, accounts for most of the growth differentials across countries. Easterly and Levine argue that “in the search for the secrets of long-run economic growth, a high priority should go to rigorously defining TFP, empirically dissecting it, and identifying the policies and institutions most conducive to its growth”. Our paper obeys their injunction by measuring the impact that trade in new and better varieties has had on productivity growth in different countries around the world.

We begin our exploration of the data by documenting that a defining characteristic of the growth of world trade over the last decade has been the import of new goods as well as existing

¹ E.g. Ethier [1982], Feenstra [1990], Jones [1995], Krugman [1980] and Rivera-Batiz and Romer [1991]

goods from new sources of supply. Analyzing 6-digit bilateral flows over the period 1994-2003, we show that had it not been for new variety growth, the share of trade to world GDP would have *fallen* in the typical developing country during this period. This provides evidence that trade in new and better products has the potential to contribute to our understanding of productivity growth around the world. Moreover we show that there are far more trade flows than would be obtained in a world of homogeneous goods. Since countries are importing new goods and these goods are differentiated, this implies gains to productivity and welfare along the lines predicted by endogenous growth models.

The next step is to quantify these gains. The starting point for our structural estimation is the work of Feenstra and Markusen [1994]. They develop a methodology for measuring the impact of new varieties on productivity. The constant elasticity of substitution (CES) specification enables us to identify the gains from variety by keeping track of only two factors: the elasticity of substitution among different varieties of a good and shifts in expenditure shares among new, remaining, and disappearing goods. The main intuition is that increasing the number of varieties does not increase productivity much if new varieties are close substitutes to existing varieties or if the share of new varieties is small relative to existing ones. We adapt this methodology and apply methods developed in Feenstra [1994] and Broda and Weinstein [2006] to obtain elasticities of substitution and supply for a large number of sectors in each of 73 countries. By letting demand and supply parameters vary at the 3-digit Harmonized-System (HS) level, we are able to divide each country's import sectors into about 200 sectors and estimate different demand and supply elasticities for varieties defined at the 6-digit HS level.

To give an example from our data, we allow for 6 different types of groundnut products: e.g. shelled raw groundnuts, unshelled raw groundnuts, prepared groundnuts, etc. Data limitations force us to treat new imports of some – but not all – groundnut products symmetrically (e.g. shelled and unshelled raw groundnuts), but nonetheless we are able to structurally estimate four different demand parameters and four different supply parameters in the groundnut market per country. These parameters enable us to estimate how new imports of unshelled groundnuts affect each economy without making any restrictions that countries value raw groundnuts similarly or that liberalizing raw groundnuts is similar to liberalizing prepared groundnuts.

We then aggregate these estimates of the gains in each sector and in each country to construct the contribution from the gains from variety on an exact productivity index. We find that the growth in new traded varieties accounts for around 0.2 percentage points per year or 30 percent of TFP growth in the typical developing country. The impact of trade on growth is smaller in developed countries, with variety growth accounting for 0.1 percentage points or 5 percent of TFP growth.

In the second part of the paper we examine the “growth effect.” Here we quantify the impact that the global rise in imported varieties in the last 10 years has on future output growth. For this purpose, we estimate the impact that size has on the creation of new varieties using disaggregate data on new traded varieties for 73 countries. We use a general dynamic specification and the data supports the semi-endogenous growth model over the purely exogenous and fully endogenous models. Thus, trade only temporarily raises growth rates over their steady state level (which is exogenous). We find the dynamic effect is small. This suggests that the ongoing increase in trade flows is an important source of rising productivity levels and thus growth rates, but the dynamic impact of these gains on future growth rates does not raise welfare much.

III. The Growth in World Varieties

Why are aggregate trade volumes rising? In order to answer this question, it will be useful to establish some terminology. We define an imported variety as a 6-digit Harmonized System (HS) product category from a particular country. To give an example drawn from our data, “raw, shelled groundnuts” is a product category, and “Gambian, raw, shelled groundnuts” is a variety. Obviously, counting categories and varieties can be problematic because new varieties may be lower or higher quality relative to existing ones. This is an issue that we will deal with in our econometric section, but it is still useful to have some sense of the changes in variety growth in the raw data.

In Table 1, we present summary statistics for the 73 countries that reported 6-digit HS import data using the 1992 classification system in the COMTRADE database for at least 5 years

between the 1994 and 2003.² Most countries have consecutive import and export data from 1994 to 2003 (see Appendix for years for each country). On average, we have approximately 3 million bilateral trade flows per year. As one can see from Table 1, the median number of imported product categories in our sample is 4193 in 1994 and 4265 in 2003.³ This indicates that the typical country in the world imports just over four fifths of all of the different types of goods (not varieties) traded internationally. In other words, import competition exists in most countries in most markets.

The change in the number of source countries per category is more striking. The median number of varieties imported by a country in our sample rose by 28 percent from 28,664 to 36,705 over the same time period. In order to assess the implications of this change for each importer, we first computed the average number of countries supplying each 6-digit import good and then computed the median across importers of these averages. In 1994, the typical country imported its goods from on average 7 suppliers. This indicates that only very few of the close to 200 potential exporters actually supplied a good into the typical importer's market. Between 1994 and 2003, however, the average number of suppliers of a good in the typical importer rose from 7.0 to just over 9.3, indicating that most countries experienced a substantial increase in the number of exporters supplying any given market. This 33 percent increase in the number of countries supplying the imports of a good to the typical country in combination with the rise in total number varieties imported by the typical country, makes a *prima facie* case that the number of varieties entering most countries rose globally.

Figures 1, 2 and 3 portrays this information graphically for our 73 countries. We plot the final number of positive import categories (in logs) against the initial level (in logs) in Figure 1. The dashed line corresponds to the 45 degree line which indicates no growth in imported goods. Almost all countries experienced an increase in the number of goods imported, and those that experienced declines, experienced only modest ones. The data indicate that there is a general increase in the number of goods imported. As one can see from the figure, countries that

² We initially had 79 countries, but we noticed some discrepancies between the COMTRADE data and the World Bank's World Development Indicators database) that made us not trust the data for a few countries. We dropped 4 countries (Costa Rica, Saint Lucia, Trinidad and Tobago, and Uganda) because the ratio of imports reported by COMTRADE relative to the WDI fluctuated by over 20 percent. We also dropped county years if the ratio of imports in the two databases differed by more than 50 percent. This eliminated Moldova and the Seychelles.

³This rise occurred despite the fact that several HS categories were retired over this time period, which caused the total number of HS categories to fall slightly from 5036 to 4980.

imported in a large set of products in 1994 had proportionally less growth in new goods than those in the rest of the sample.

This inverse relationship between growth in new goods and the initial level of goods with positive imports suggests two potential limitations of the data. First, even at the 6-digit level, most countries import most goods (and some countries import virtually all goods), so the scope for identifying variety growth through the new goods dimension is going to be limited. This is especially true for large countries that tend to import in virtually all categories. Second, the total number of existing product categories is bounded above. Statistical offices define goods categories based on the existing set of goods, not future goods. This means that if variety growth manifests itself in new goods, new varieties will initially be placed in the same categories as existing goods, and we will underestimate variety growth.

Figure 2 shows that the growth in the number of source countries per good is also larger in countries that imported from fewer sources per good in 1994. As one can see from Figure 2, virtually all countries in the world shared this common trend of importing from more countries. Figure 3 reports the pattern of growth depending on the initial level of varieties. It shows that variety growth is positive for almost all countries in our sample. The only outliers in these data are countries with some major macroeconomic or political disruption such as Argentina, Central African Republic, Malaysia and Uruguay. Moreover, this figure shows that import variety growth was a common feature of all countries in the sample during this period, despite the fact that some countries could not import in more categories.

Table 2 documents the pattern of growth in trade for our sample of countries. First, we compute the growth rate of imports as a share of GDP for each importer in our sample. The fact that world trade has been growing at a faster pace than world GDP has been documented in numerous studies, and this is a feature of our data as well. For the median developed country, the ratio of imports to GDP grew 29 percent over our sample period and around 16 percent for the median developing country. The novelty of using disaggregate trade data is that we can account for how much of this growth was due to the growth in imports of the varieties that were already present in 1994 and how much was due to the appearance of new varieties. Interestingly, we find that for the typical developing country, had it not been for the new varieties trade as a share of GDP would have *fallen* during this period. This underscores the importance that new and better

products have in the developing world. For the median developed country, we find that half of the growth in imports as a share of GDP comes from the import of new varieties.

A similar finding concerns the sources of the growth in imports around the world.

Because average growth rates can be heavily influenced by countries with low initial import levels, we report median import growth rates and the medians of the constituent growth rates in Table 2. The growth rate of varieties imported in 1994 to GDP was only 15 percent for developed countries and *negative* 16 percent for developing countries. This means that about half of the growth in import GDP ratios was due to the import of new goods in the typical developed country. Had the typical developing country not imported new varieties, its trade to GDP ratio would have fallen. Of course, we can't know what would have happened to the import levels of initial imports if new varieties had not appeared, but the data makes a strong case for thinking about the import of new varieties when thinking about why trade to GDP ratios are rising

The growth in the number of source countries supplying any good is indicative of a rise in the level of global integration: most countries already import 10-ton trucks, but now they do so from more sources of supply. To get some sense of how far away the world is from the fully integrated equilibrium, we can think of two useful benchmarks. The first is minimal trade dispersion. In this benchmark, conditional on a variety being exported, it is imported by only one country. The second is maximal trade dispersion: conditional on being exported, a variety is imported everywhere. One simple way of measuring how close we are to either extreme is to divide the total number of varieties imported by all countries by the total number of varieties that would be imported if every exported variety were imported everywhere. Since, by definition, every good that is exported must be imported by some country, the ratio of actual varieties imported to potential varieties in an industry will range from one over the number of possible importers (1/72) to one. We therefore normalized this index as follows so that 0 corresponds to minimal trade dispersion (each exported good going to only one destination) and 1 to maximal dispersion (each exported good imported by every country):

$$(1) \quad \text{Trade Dispersion} = \left(\frac{\# \text{ of Imported Varieties}}{\# \text{ of Exported 6-digit Categories} * 72} - \frac{1}{72} \right) * \frac{72}{71}$$

We present the results from the trade dispersion index in Table 3. Overall, this index grew by almost 20 percent between 1994 and 2003. The number of exported varieties worldwide

only grew by 3 percent, which means that most of the growth comes from the increased number of countries importing any given variety.

Despite this rapid increase in varieties imported even relative to the increase in available varieties, our “dispersion index” stood at just above 15 percent in 2003. In other words, conditional on a good being exported, only about one in seven of the world’s countries import it. This tells us two important facts about globalization. First, there has been a substantial rise in the number of countries importing any internationally traded good. Second, we are very far away from a world in which an exported good is imported by anything more than a small fraction of potential importing countries.

The growth in these new varieties would have little meaning if imports of the same good from different countries are perfect substitutes. One way to confirm whether goods at this level of aggregation are indeed differentiated is by testing a simple prediction by Dorfman, Samuelson and Solow [1958]. They argued that in a world in which goods are perfectly homogeneous and firms are perfectly competitive, one can model trade as a linear programming problem in which one minimizes trade costs subject to the constraint that one must satisfy every country’s net offer of each good. With C countries in the world, Dorfman, Samuelson, and Solow show that the solution to any such minimization problem will entail no more than $C - 1$ positive bilateral trade flows for any good. The intuition for this result is that if goods are homogeneous and trade costs are positive, countries will not be on both sides of the market, i.e. they will not export and import the same good. Thus, the most flows will obtain when there is one exporter and $C - 1$ importers. If there are more exporters of a good, then trade between them will be zero and the number of flows will tend to fall or at most remain the same.

The “ $C - 1$ ” condition provides a simple test for assessing what share of world trade could plausibly be modeled as the exchange of homogeneous goods. In order to examine this, we analyzed the trade flows for each good in which there exists some trade among the 73 countries that reported trade data in 2003. In each of the 5036 categories in which at least one country in our sample exported to another country in a 6-digit category, we counted the number of positive bilateral trade flows. The frequency distribution is presented in Figure 4. The line in the figure separates those sectors satisfying the necessary condition for being a homogeneous good, i.e. having 72 or fewer flows. Only 4 percent of the goods at the 6-digit satisfy this condition. This suggests that even at the 6-digit level, it would be a mistake to model more than a small fraction

of world trade as trade in homogeneous goods. However, if the vast majority of trade is in differentiated products and imported varieties are rising, then this implies that there may be gains from importing new varieties.

To sum, our examination of the raw data reveals several important features of globalization. First, in most countries, the trade to GDP ratio is rising not because they import more of the same varieties, but rather because they import new goods and, in particular, new varieties of existing goods. This phenomenon is critical for understanding the growth of trade relative to GDP around the world. Second, there appears to be vastly more bilateral flows between countries than one would have predicted if all 6-digit goods were homogeneous. This implies that countries may have gains from the import of new varieties. How much the increase in new varieties matters for global growth is the issue we address in the next section.

IV. Empirical Strategy

A. Theory

There are many versions of endogenous growth models that we could use as the basis of our estimation. Here we will focus on the ones in which growth in varieties play an important role in economic growth (e.g. Jones [1995] and Rivera-Batiz and Romer (RBR) [1991]). Following the literature, we term this the “level effect” and will focus our empirical estimation on the class of models that specifies the love of variety production functions in terms of a CES function. These models differ in important ways in their specification of innovation or the “growth effect”, and we will address these differences in a second section⁴

A.1. The Level Effect

We start with a specification of the production technology where there is only one final output produced with labor and a set of intermediate or capital goods indexed by g .⁵ Technological progress is represented by the creation of designs for the new types of intermediate or capital inputs. In particular, let N be the most recently invented design, A be a standard productivity parameter and output take the following form

⁴ These models can also be interpreted to be models of how new ideas (as opposed to varieties of goods) affect growth. If these channels are important then the forces identified by these models would be even more important.

⁵ Throughout the empirical work we will not distinguish capital from intermediate products

$$(2) \quad Y_t = (AL_t)^\alpha \left(\sum_1^N x_{gt}^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta(1-\alpha)}{\eta-1}} ; \eta > 1$$

where α is the share of labor in output, and η is the elasticity of substitution between varieties of goods.

In a static setup, the effect that trade has on output can be simply traced by the effect that trade has on the number of goods, N , that are used in production. In particular, under standard preference and monopolistic competition assumptions (like in Krugman [1980] or Grossman and Helpman [1990] and Rivera-Batiz and Romer [1992]), it can be shown that $x_{gt} = \bar{x}_t \forall g$. In this case, when two symmetric economies fully integrate, the number of intermediate products jumps from N to $2N$. This generates a “level effect” as each country’s GDP will increase because each enjoys more product variety through trade. It is straightforward to decompose the output change into changes in A , N and factors of production. In particular,

$$(3) \quad \Delta \ln Y = \Delta \ln \tilde{A} + \alpha \Delta \ln L + (1 - \alpha) \Delta \ln \bar{x}$$

$$(4) \quad \Delta \ln \tilde{A} = \Delta \ln A + \frac{(1 - \alpha)}{\eta - 1} \Delta \ln N$$

where \tilde{A} is a composite TFP measure that includes advances in new products and all other sources of TFP summarized in A . As is apparent from equation (4), the impact of an increase in the number of intermediate or capital goods is larger, the higher is the share of these goods in final output (higher $(1 - \alpha)$), and the less substitutable are new varieties with existing ones (higher η). In this simple symmetric world, it can be easily shown that

$$\Delta \ln Y = \Delta \ln \tilde{A} = \frac{(1 - \alpha)}{\eta - 1} \ln 2.$$

In other words, trade increases GDP as the rise in the number of imported goods increases productivity.

A.2. The Growth Effect

A central component of endogenous growth models is, quite naturally, the innovation function, i.e. the technology with which the economy produces new and better products. In the

knowledge-driven specification of innovation, new products are generated by the labor force dedicated to research, L_A , according to the following equation:

$$(5) \quad \dot{\tilde{A}} = \delta L_A^\xi \tilde{A}^\phi .$$

where $\delta > 0$, $0 < \xi \leq 1$ and $\phi \leq 1$. If ξ and ϕ equal one, then this reduces to the Romer version of the lab-equipment specification of the endogenous growth model. Jones [1995] considers the case in which both of these parameters are less than unity. Since this specification postulates that innovation is only a function of labor and the stock of knowledge, there is no channel through which imports can raise TFP growth.

An alternative approach is the “lab-equipment” specification in which innovation requires capital or intermediate inputs. In this case new imported intermediates will raise innovation because intermediates lower the cost of doing research. Eicher and Turnovsky [1999] extend the endogenous and semi-endogenous models to this more general setting which allows for intermediates to matter for growth and derive the balanced growth path. Although we won’t consider the most general possible version of the lab-equipment model, we can estimate parameters for a case in which $\dot{\tilde{A}}$ is allowed to increase in Y and A . In this class of models, the innovation equation can be written as

$$(6) \quad \dot{\tilde{A}} = \delta F^\chi \tilde{A}^\phi ,$$

where $0 < \chi \leq 1$ and $\phi \leq 1$, $F = L^{2/3} K^{1/3}$, K is the capital stock, and the exponents correspond to the labor and capital shares.⁶

If χ and ϕ are less than unity, then this specification will result in a temporary effect of trade openness on growth. In particular, if the level effect raises F because of the love of the variety production process, then equation (6) indicates that this will raise the rate of innovation. In the steady state, however, growth will be unaffected as in Jones [1995]. Nevertheless, there will be a temporary impact of imports on growth.

B. Implementation

In order to move away from the stylized world presented in section A.1 and closer to the empirical specification that we use, we introduce a number of asymmetries in the production

⁶ Here, we assume a value of the share of capital in output of 1/3, which is the same used by Hall and Jones [1999] and is broadly consistent with national income accounts data for developed countries.

function. We start by identifying the goods that are produced domestically from those that are imported. Without loss of generality, assume that goods 1 to J_t are produced domestically, and goods $J_t + 1$ to N_t are imported. This allows us to trace the effect on GDP of imported and domestic goods separately. We can think of $x_{gt} \forall g \in [J_t + 1, \dots, N_t]$ as a composite of differentiated varieties of a particular imported input. A convenient way to aggregate over varieties of each $g \in [J_t + 1, \dots, N_t]$ is to use a CES function:

$$(7) \quad x_{gt} = \left(\sum_{v \in I_{gt}} d_{gvt} x_{gvt}^{\frac{\sigma_g - 1}{\sigma_g}} \right)^{\frac{\sigma_g}{\sigma_g - 1}} ; \sigma_g > 1 \quad \forall g \in [J_t + 1, \dots, N_t]$$

where I_{gt} is the set of varieties with positive imports in time t , i.e. $x_{gvt} > 0$. Notice that the elasticities of substitution between varieties are allowed to vary across imported goods g . For future reference, we define $I_g = I_{gt} \cap I_{gt-1}$ as the set of varieties that are common over time for good g , $I_t = \bigcup_{g=J_t+1}^{N_t} I_{gt}$ as the set of all available imported varieties in period t , and $I = I_t \cap I_{t-1}$ as the set of the varieties that are common over time for all goods.

With this structure we can apply the same methodology as in Feenstra and Markusen [1994] to decompose output growth into higher quantity of existing inputs and greater range of intermediate imported inputs. We begin by assuming that the number of domestic goods is unchanged, as in Krugman [1980] and Romer [1990], and focus on the impact of new varieties on GDP. Later in this section, we derive the impact that increased GDP has on the number of domestically produced intermediate inputs. By using the relationship between output and minimum unit-costs, $Y_t = E_t / c(p_t, I_t)$, where c is the minimum unit cost function of the production function Y_t . Given the assumptions underlying equation (2), we can decompose output changes into the contribution of new and existing inputs:

$$(8) \quad \begin{aligned} \ln(y_1 / y_0) &= \ln(E_1 / E_0) - \ln \left[c(p_1, I_1) / c(p_0, I_0) \right] \\ &= \ln(E_1 / E_0) - \ln \left[c(p_1, I) / c(p_0, I) \right] - \ln \left[c(p_1, I_1) / c(p_1, I_0) \right]^2 \end{aligned}$$

where the last term of equation (8) traces the change in the minimum unit cost function that is due to an increase in variety from I_0 to I_1 . Following Feenstra [1994] and Broda and Weinstein

[2006], we show that given equations (2) and (7), the ratio of unit costs in the first line of equation (8) takes the following simple expression:

$$(9) \quad \frac{c(p_1, I_1)}{c(p_0, I_0)} = \frac{c(p_1, I)}{c(p_0, I)} \times \prod_{g=J_t+1}^{N_t} \left(\frac{\lambda_{gt}}{\lambda_{gt-1}} \right)^{\frac{w_{gt}}{\sigma_g - 1}},$$

where $\lambda_{gt} = \frac{\sum_{v \in I_g} p_{gvt} x_{gvt}}{\sum_{v \in I_{gt}} p_{gvt} x_{gvt}}$ and $\lambda_{gt-1} = \frac{\sum_{v \in I_g} p_{gvt-1} x_{gvt-1}}{\sum_{v \in I_{gt-1}} p_{gvt-1} x_{gvt-1}}$ measure the extent of variety in each

period, and w_{gt} are Sato-Vartia log-ideal weights:

$$(10) \quad w_{gvt} = \frac{\frac{s_{gvt} - s_{gvt-1}}{\ln s_{gvt} - \ln s_{gvt-1}}}{\sum_{v \in I_g} \left(\frac{s_{gvt} - s_{gvt-1}}{\ln s_{gvt} - \ln s_{gvt-1}} \right)}, \text{ where } s_{gvt} = \frac{p_{gvt} x_{gvt}}{\sum_{v \in I_g} p_{gvt} x_{gvt}}.$$

Equation (9) states that the exact price index with variety change is equal to the “conventional” price index, $c(p_1, I_0)/c(p_0, I_0)$ (i.e., the exact price index of the common varieties over time), multiplied by an additional term which captures the role of the new and disappearing varieties.⁷ The last term in equation (9) tells us the gain in the level of GDP that can be directly traced to increases in imported variety from $(N_0 - J_0)$ to $(N_1 - J_1)$. In terms of the decomposition of TFP in (4), the log of the geometric weighted average of λ 's in (9) replaces the simple term that includes the change in N , to give the following expression:

$$(11) \quad \Delta \ln \tilde{A} = \Delta \ln A - \sum_{g=J_t+1}^{N_t} \frac{w_{gt}}{\sigma_g - 1} \ln \left(\frac{\lambda_{gt}}{\lambda_{gt-1}} \right).$$

In the next sections, we will describe how to empirically estimate the relevant parameters to calculate the share of TFP changes that are due to trade in new varieties. Note that λ_{gt} equals the fraction of expenditure in time t on the varieties that are available in both periods (i.e., $v \in I_g = (I_{gt} \cap I_{gt-1})$) relative to the entire set of varieties available in period t (i.e., $v \in I_{gt}$). Thus, this new term implies that the higher the expenditure share of new varieties, the lower is λ_{gt} , and the smaller is the exact price index relative to the conventional price index. In the fully

⁷ All of the index numbers used in this paper suffer from the classic “index number problem”. In particular, results are dependent on the base year or years used. Since we are examining long-run changes, we use two base years 1972 and 1990.

symmetric case (i.e., $\sigma_g = \eta \forall g$), the change in log GDP due to new inputs we would obtain from equation (9) is exactly equal to $(1 - \alpha) \ln 2 / (\eta - 1)$, where the 2 reflects the doubling in the number of available varieties when identical countries trade freely. It is easy to see that, in this case, an increase in the number of varieties leads to a fall in the exact price index and an increase in GDP.

The Feenstra price index also depends on the good-specific elasticity of substitution, σ_g . As σ_g grows, the term $1/(\sigma_g - 1)$ approaches zero, and the bias term $(\lambda_{gt} / \lambda_{gt-1})^{1/(\sigma_g - 1)}$ becomes unity. That is, when existing varieties are close substitutes to new or disappearing varieties changes in variety will not have a large effect on the exact price index. By contrast, when σ_g is close to unity, varieties are not close substitutes, $1/(\sigma_g - 1)$ is high, and therefore new varieties are very valuable, and disappearing varieties are very costly.

V. Econometrics

We rely closely on the methodology derived in Feenstra [1994] as extended by Broda and Weinstein [2006] to estimate elasticities of substitution between varieties of imported goods. As opposed to in the previous section, we now index each variable with a country subscript i to emphasize that elasticities are estimated separately for each good and importing country. We estimate the following system of import demand and export supply equations:

$$(12) \quad \Delta^{k_{ig}} \ln s_{igvt} = -(\sigma_{ig} - 1) \Delta^{k_{ig}} \ln p_{igvt} + \Delta^{k_{ig}} \varepsilon_{igvt}$$

$$(13) \quad \Delta^{k_{ig}} \ln p_{igvt} = \frac{\omega_{ig}}{1 + \omega_{ig}} \Delta^{k_{ig}} \ln s_{igvt} + \Delta^{k_{ig}} \delta_{igvt}$$

where $\Delta^{k_{ig}} x_{igvt} = \Delta x_{igvt} - \Delta x_{igk_{ig}t}$ (i.e., differencing across two different varieties of a given i - g pair), i denotes the importer country, g a 4-digit good, and v (for variety) a particular variety of good g , and $s_{igvt} = p_{igvt} x_{igvt} \cdot \varepsilon_{igvt}$ are taste or quality shocks to variety v of good g in country i and δ_{igvt} are shocks to the supply of the same variety.

Equation (12) can be thought of as the optimal demand for intermediate varieties of good g derived from a CES final good production function, and (13) is the supply of that variety expressed in terms of shares. In particular, the inverse elasticity of supply is given by ω_{ig} which is allowed to be different from zero but restricted to be the same for all varieties within an i - g

pair. More importantly for the identification strategy is our assumption that

$E\left(\Delta^{k_{ig}} \varepsilon_{igvt} \Delta^{k_{ig}} \delta_{igvt}\right) = 0$. That is, once good-time specific effects are controlled for, demand and supply errors at the variety level are assumed to be uncorrelated.

To derive the key moment conditions that will be used for identification, it is convenient to multiply (12) and (13) together to take advantage of the independence condition of errors:

$$(14) \quad \left(\Delta^{k_{ig}} \ln p_{igvt}\right)^2 = \theta_{11} \left(\Delta^{k_{ig}} \ln s_{igvt}\right)^2 + \theta_{12} \left(\Delta^{k_{ig}} \ln p_{igvt} \Delta^{k_{ig}} \ln s_{igvt}\right) + u_{igvt}$$

where

$$\theta_{ig1} = \frac{\omega_{ig}}{(1 + \omega_{ig})(\sigma_{ig} - 1)}, \quad \theta_{ig2} = \frac{1 - \omega_{ig}(\sigma_{ig} - 2)}{(1 + \omega_{ig})(\sigma_{ig} - 1)} \quad \text{and} \quad u_{igvt} = \Delta^{k_{ig}} \varepsilon_{igvt} \Delta^{k_{ig}} \delta_{igvt}.$$

Unfortunately, $\beta_{ig} = \begin{pmatrix} \theta_{1ig} \\ \theta_{2ig} \end{pmatrix}$ cannot be consistently estimated from (15) as the error term, u_{igvt} , is

correlated with the regressands that depend on prices and expenditure shares. However, it is still possible to obtain consistency by exploiting the panel nature of the dataset combined with the assumption that demand and supply elasticities are constant over varieties of the same good. In particular, we can define a set of moment conditions for each good g and each importing country i , by using the independence of the unobserved demand and supply disturbances for each variety over time, i.e

$$(15) \quad G_v(\beta_{ig}) = E_t\left(u_{igvt}(\beta_{ig})\right) = 0 \quad \forall v, g \text{ and } i.$$

For each good g and importer i , all the moment conditions that enter the GMM objective function can be stacked and combined to obtain Hansen's [1982] estimator:

$$(16) \quad \hat{\beta}_{ig} = \arg \min_{\beta_{ig} \in B} G^*(\beta_{ig})' W G^*(\beta_{ig}) \quad \forall g \text{ and } i.$$

where $G^*(\beta_{ig})$ is the sample analog of $G_v(\beta_{ig})$ stacked over all varieties v of a good g , W is a positive definite weighting matrix to be defined below, and B is the set of economically feasible parameters β_{ig} , which is common across importers and goods (i.e. $\sigma_{ig} > 1$ and $\omega_{ig} > 0 \forall i, g$). We follow Broda and Weinstein [2006] in the way we implement this optimization. We first estimate the “between” version of (14) to obtain estimates of θ_{11} and θ_{12} and then solve for β_{ig} as in Feenstra [1994]. If this produces imaginary estimates or estimates of the wrong sign we use a

grid search of β s over the space defined by B . In particular, we evaluate the GMM objective function for values of $\sigma_{ig} > 1$ at intervals that are approximately 5 percent apart.⁸

The problem of measurement error in unit values motivates our weighting scheme. In particular, there is good reason to believe that unit values calculated based on large volumes are much better measured than those based on small volumes of imports. In the appendix of Broda and Weinstein [2006], they show that this requires us to add one additional term inversely related to the quantity of imports from the country and weight the data so that the variances are more sensitive to price movements based on large shipments than small ones. The use of the between estimate coupled with our need to estimate σ_{ig} , ω_{ig} , and a constant means that we need data from at least three exporting countries for each importer in each good and at least three two time differences to identify β .

VI. Results

We begin by characterizing the growth in world varieties in terms of our key parameters. One of the problems that we face is the sheer number of goods and countries in our dataset. With 73 countries and most countries reporting close to 200 3-digit HS sectors, we constructed approximately 13,000 lambda ratios and estimated an equal number of elasticities of substitution. It is obviously impossible to report all of these and far too much to hope that there won't be some outliers. One approach to assessing the reasonableness of these estimates is to compare them with existing estimates and our priors.

Broda and Weinstein [2006] used the same methodology to estimate elasticities at various levels of aggregation. Our study differs from that one not only in scope but also in the fact that we have moved from 10-digit to 6-digit data. The first question we have is how much is lost from this switch. Since the US is one of our 73 countries, we can compare the estimates of Broda

⁸ For computational easiness, we performed the grid search over values of σ_{ig} and ρ_{ig} where ρ_{ig} is related to ω_{ig} in the following way: $\omega_{ig} = \rho_{ig} / [\sigma_{ig} (1 - \rho_{ig}) - 1]$. The objective function was evaluated at values for $\sigma_{ig} \in [1.05, 131.5]$ at intervals that are 5 percent apart, and for $\rho_{ig} \in [0.01, 1]$ at intervals 0.01 apart. Only combinations of σ_{ig} and ρ_{ig} that imply $\sigma_{ig} > 1$ and $\omega_{ig} > 0$ are used. To ensure we used a sufficiently tight grid, we cross-checked these grid-searched parameters with estimates obtained by non-linear least squares as well as those obtained through Feenstra's original methodology. Using our grid spacing, the difference between the parameters estimated using Feenstra's methodology and ours differed only by a few percent for those σ_{ig} and ω_{ig} for which we could apply Feenstra's "between" approach.

and Weinstein [2006] with the ones we obtain here. Unfortunately, we cannot do this precisely because the earlier study computed sigmas and lambdas for the years 1990-2001 while we are using the years 1994 to 2003. Even so, our median sigma for the US is 2.3 as compared with the 3-digit median of 2.5 obtained in Broda and Weinstein. The closeness of these estimates indicates that the switch in time periods and the higher degree of aggregation is not producing major movements in the typical parameter estimates for the US.

Given that our data set incorporates countries with varying degrees of development, it is reasonable to worry about whether our methodology works for the typical country. In order to provide a sense of the distribution of elasticities of substitution, we computed the average and median values of each parameter for each country and then compute the distribution of these parameters across countries. The typical country has a median sigma of 3.4, significantly larger than that of the United States. Average sigmas tend to be higher than medians because the sigmas are bounded below by 1. After sorting countries by their average sigma, we find that the typical country has an average sigma of 6.8, while the US has an average sigma of 4.2, suggesting that on average the US tends to value variety somewhat more than the typical country. Another way of looking at the results is to focus on the results for the sub-samples of developed and developing countries. The bottom of Table 4 reveals that the median elasticity of substitution does not vary by the level of development. This indicates that there is no strong relationship between income per capita and the elasticity of substitution across countries. Not surprisingly, there are some outliers. The United States and Greece have somewhat low sigmas and Sweden and Canada have high sigmas. Overall the median elasticity of substitution shows a fair bit of dispersion, with the minimum median elasticity being 2.3, while the maximum is more than twice as large.

This, of course, raises the question of whether our elasticity estimates themselves are sensible. One approach to assessing their “reasonableness” is by comparing them with our priors. One might suspect that varieties of goods traded on organized exchanges are more substitutable than those that are not. For example, natural gas exported by different countries is likely to be more similar than telecommunications equipment emanating from different suppliers. Rauch [1999] classifies all 4-digit SITC product categories in 3 groups: those sold in organized exchanges, those that have a reference price in the US and the rest. Broadly speaking, the classification helps distinguishing between products that are commodities and those that are

differentiated. We obtain concordances between 3-digit HS codes and 4-digit SITC codes to group our estimates of elasticities of substitution into the groups distinguished by Rauch. Table 5 shows median and mean according to these different groups. In all cases, we can strongly reject the hypothesis that the median or mean for the group of commodity products is lower than that for the other two groups. In particular, the average elasticity is 12.1 for commodities while it is around 7.2 for the rest of the products. Thus, our elasticity estimates seem to be plausible by this criterion.

Our estimates of the gains from new varieties depend on two factors: how differentiated varieties are and the importance of new imported varieties. While the elasticity estimates give us information about the former, the lambda ratios provide information about the importance of net variety creation in any given market. If we sort countries by the median lambda ratio in each of their import sectors, we find that the median lambda ratio in the typical country is 0.93. This suggests that the typical country experienced a net increase in varieties (creation less destruction) of 7.1 percent over 10 years in the typical sector or about 0.7 percent per year.

There is, of course, substantial variation in median lambda ratios across countries, as one can see in the last column of Table 4. Some countries have seen the effective number of varieties in their typical import sector almost double. However, there is an unmistakable pattern in the data. All but 3 countries experienced an increase in variety in the typical industry, indicating that the increase in varieties is a global phenomenon. The only countries that experienced a fall in variety in the typical industry (median lambda ratio higher than 1) were Argentina, Sri Lanka and Uruguay. The experiences of Argentina and Uruguay may reflect major economic crisis that rocked these countries in 2002 and 2003, respectively, while Sri Lanka's experience may be related to its civil war. On the other hand, several countries that liberalized extensively over this period, such as Poland and India, experienced fairly substantial declines in their median lambda ratios, indicating that the variety of imports in their representative industries rose substantially. Similarly, some of the EU countries in our sample had lower lambda ratios than the median indicating fairly substantial gains in variety. Whether these results reflect the impact of the European trade liberalization, the Polish opening to international trade, and Indian liberalization is difficult to say because we do not know what liberalizations occurred in countries without substantial drops in their lambda ratios, but the results are clearly suggestive of such a link.

Using the lambda ratios and elasticities of substitution for each good in each country we obtain an estimate of the productivity gains from new varieties (see equation (15)). The distribution of the TFP gains calculated on a per-year basis is reported in Table 4. The growth in new varieties over the period 1994-2003 increased productivity by 0.13 percent per year for the typical country in our sample. Some developed countries, like Australia, Denmark, Ireland, New Zealand, Spain, The Netherlands, The United Kingdom also benefited substantially, but most of the productivity growth in many of the largest countries cannot be accounted for by new imports. Some of this may be due to insufficient aggregation. Large economies typically import most goods at the 6-digit level and hence we cannot measure the gains from many new varieties which appear in categories in which the country already imports. Indeed, the US, which has the second smallest gain among developed countries from imported variety in the 6-digit data, had a gain that was 3 times larger when Broda and Weinstein [2006] computed the same value using 10-digit data for the period 1990-2001.⁹ Almost surely, most of this difference is due to the fact that there are only about one third as many 6-digit categories as 10-digit categories. Indeed, if we drop the ten largest countries in the world (as measured by GDP), the median country's productivity gain due to new imported varieties rises from 0.13 to 0.15 and the median for developed rises by over 30 percent to 0.11. This suggests that we may have a significant downward bias in our estimates of the gains from new varieties for large countries.

The average and median impacts of new varieties on productivity growth are presented in the first two rows of Table 7. On average, new imported varieties raised productivity growth by 0.27 percent per year and the median impact was 0.13 percent per year. These numbers are both significantly different from zero. If we restrict ourselves to the sample of countries for which we can compute TFP from the Penn World Tables v6.1 data, we find that productivity gains from new imported varieties are 15 percent as large as the average country's per capita growth rate or TFP growth rate and 12 percent the size of the median value. This indicates that while there are other important factors that determine TFP growth, the import of new varieties has a substantial impact.

As one might have expected in a diverse sample of countries, there is a lot of variance in individual country experiences. The median developed country's productivity growth was about

⁹ We also estimated the level effect for Japan using HS 9-digit data over the period 1992 to 2005 and similarly found that the level effect was 3 times larger.

2 percent per year, but the median contribution of imported variety growth to productivity was only 0.1 percent per year, suggesting that for the typical developed country, new imported varieties are only a small part of the story behind their productivity growth. The impact of new varieties on developing countries is substantially higher. The typical developing country saw its productivity rise by 0.13 to 0.17 percent per year (depending on the sample) due to new imported varieties. If we use the sample of countries that also appear in the Penn World Tables, we find that the median developing country's productivity gain due to new imported varieties is 22 percent as large as the typical developing country's productivity growth. This suggests that new variety growth can account for about a quarter of developing country productivity growth over this time period. Indeed the correlation between TFP growth and our measure of productivity gains due to new varieties is a statistically significant 0.3, indicating a positive association between these two variables. Obviously, there are many forces other than new imported varieties that affect TFP growth, but our estimates suggest that the gains from this channel are not negligible.

VII. Why are varieties rising? The Growth Effect

Thus far, we have been concerned with documenting that the import of new varieties is growing, that varieties are differentiated, and that this process has been exerting a non-trivial effect on productivity growth in developing countries. Endogenous growth models postulate the existence of a “growth effect” that arises from a linkage between the rate of new variety creation and research activity in each country. Although the parameters linking the variables differ across models, all models in the Romer tradition predict that as countries become larger and more productive, the rate of variety creation will rise. In this section we first test whether this is a feature of the data, and then use our estimated parameters to estimate the growth effect.

In order to do this, we need to obtain a good measure of the change in varieties.¹⁰ Ideally, we would have counts of all new varieties produced in a country. Unfortunately, we have no information about the number of non-traded varieties produced in any of our countries. However, if we make some reasonable assumptions, we can identify associations between observables and new exported varieties. There are two polar ways to link country observables with new exported

¹⁰ There is, however, a vast literature on the determinants of innovation at the firm level (see Klette and Kortum (2004) for a summary of the stylized facts of innovation at the firm level)

varieties. One is to assume that all new products are exportable. In this case, the relevant measure of F is the national value, and we can specify the right-hand side of equation (6) as it is written. A second approach is to assume that the amount of R&D in the tradable goods sector is equal to the fraction of the economy that is devoted to exports. In this “export specification” case, we should multiply F by the share of exports in the economy.

A second problem that we face implementing our test is that some countries in our sample export virtually every 6-digit good. Figure 5 plots the number of sectors with no exports against the level of exports in 1994. Clearly, the largest exporters already export in almost every product category and cannot substantially increase the number of new goods they export. For example, the US and Germany had positive exports in all but a couple of the 5036 possible sectors. This means that it is almost impossible for us to measure the value of innovation in these economies.

Presumably, the reason why we cannot measure new variety creation in the US is not because the possibility of innovation ended in 1994, but rather because we cannot measure the creation of a new good if it gets classified in a sector in which the US already exports. For example, many new technologies such as LCD monitors or laptop computers did not warrant their own categories at the start of our sample, but constitute new products nonetheless. We can get around this problem if we assume that innovation is equally likely to occur in any of the sectors but we can only measure innovation in a sector when a country is not exporting in a sector initially. In this case, for a common rate of innovation, we would expect to count more new exports emanating from a country that is exporting in very few sectors initially than in one that is exporting in many sectors. At the extremes, we would be able to observe every innovation in a country with no initial exports and no innovation in a country that exports in every category. In statistical terms, we can think of the number of sectors with no exports initially as the “exposure” of the country, i.e. how many possible times we can measure new goods.

Theory dictates that we should measure \hat{A} in equation (6) by the count of new exported varieties in a country over the sample period. This implies that we can estimate equation (6) using a negative binomial regression model.¹¹ The structural interpretation of the offset term in the negative binomial is that it corresponds to the differences in R&D quality in our sample of

¹¹ A key assumption in standard Poisson regression models is that the variance of the distribution of counts equals the mean. In our sample, we have the common problem of over-dispersion, where the variance is higher than the mean. This implies that the negative binomial distribution is more appropriate for our application.

countries. Formally, we assume that variation in the quality of an R&D dollar spent in different countries can be modeled as a random effect, μ_i . It is useful to define a variable, κ_i , as the the right-hand side whatever innovation equation we are estimating. For example, if we were estimating equation (6), we would set

$$(17) \quad \begin{aligned} \dot{\tilde{A}}_i &= Poisson(\kappa_i) \\ \kappa_i &= E_i \exp(\chi \ln F_i + \phi \ln \tilde{A}_i + \mu_i) \\ \exp(\mu_i) &= Gamma\left(\frac{1}{a}, a\right) \end{aligned}$$

where we make an adjustment for the exposure, E_i , and a is constant that governs the variance of R&D quality.

Table 8 presents the estimates of equation (17). Our estimates of χ are significantly greater than zero but smaller than one, which is predicted by the endogenous growth models but not the exogenous growth model. Clearly, our data supports the assumption that countries that larger countries have higher rates of innovation. A more interesting result concerns the coefficient on TFP, ϕ . This parameter is significantly greater than zero but less than one: a result that is only predicted by the Jones (1995) model of semi-endogenous growth. Taken together, the results indicate that new variety creation rates are linked to endowment and productivity levels in a manner consistent with the semi-endogenous growth model.

We now turn to estimating the economic significance. Our parameter estimates indicate that the semi-endogenous growth model is the correct way of thinking about the data. In this model, the impact of the level effect on future growth tapers off over time and does not change the steady-state growth rate. Jones [1995] derives the following relationship between per capita growth rates, g , and the variables of interest around the steady state:

$$(18) \quad g = \alpha(x - x^*) + (1 - \alpha)i(z - z^*) + g^*$$

where x is the TFP growth rate, z is the output to capital ratio, variables with asterisks correspond to steady-state values, and i is the investment to GDP ratio. The first term in equation (18) captures the fact that a shock to the TFP growth rate raises per capita income by raising labor productivity and the second term captures impact of deviations in the output to capital ratio. Movements in x and z are governed by the following laws of motion:

$$(19) \quad \frac{\dot{x}}{x} = -(1 - \phi)(x - x^*)$$

and

$$(20) \quad \frac{\dot{z}}{z} = \alpha(x - x^*) - \alpha i(z - z^*).$$

Jones considered a case in which there was a permanent increase in R&D. Here, we will consider a case in which there is a permanent increase in GDP due to imported imports. The greater variety of imports raises F in equation (6) by the amount of the level effect. This means that the growth rate of TFP, $(x - x^*)$, would rise by χ times the level effect, where we obtain our estimate of χ from Table 7. Similarly, the gains in variety imply that output per unit of capital will rise by an amount equal to the productivity gain.¹²

We investigate three methods of evaluating the impact of new varieties on growth. First, we calibrate the fully endogenous model by setting $\phi = 1$. Second, we set ϕ equal to the estimated value from Table 7 in the output specification (0.69) and finally we set it equal to the estimated value from the export specification (0.3). We choose the same values of the remaining parameters as those chosen by Jones [1995]: $g^* = x^* = 0.02$, $i = 0.3$, and $\alpha = 2/3$. If we take the median impact of new imported varieties on productivity of 0.13 percent per year, then this implies that the median country's productivity rose by 1.6 percent due to new imported varieties between 1994 and 2003. This implies that $(x - x^*) = 0.003$ and $(z - z^*) = 0.035$ initially. If we plug these values into equations (18) through (20), we can compute the dynamic path of per capita income.

The results from this exercise are plotted in Figure 6. In the fully endogenous specification, the productivity gain associated with the level effect is associated with a permanent increase in the growth rate of 0.12 percent per year. This reflects the fact that larger economies grow faster in this model. The data, however, support the semi-endogenous growth models over this specification, so we will focus on the results obtained from our estimates of ϕ rather than the assumed value in the Romer model. The growth responses for the export and output specifications are presented in Figure 6 as well. Both of these specifications yield similar predictions about the growth effect, but in these specifications, the growth effect steadily decays to zero rather than permanently raising long-run growth. The half life of the growth effect on per capita output is 13 years in both cases.

¹² We assume that the growth rate of output rises by the amount of the level effect and that z^* rises from 0.3, the level Jones used to calibrate his model, to this level multiplied by one plus the level effect (0.3048).

The economic significance of the growth effect on income can also be computed. We first compute the path of future per capita GDP relative to the current level that obtains as a result of the dynamic impact of the level effect based on the previous exercise. If we discount these future gains by 5 percent per year and subtract from them the present discounted value of future gains if the economy had only grown at the steady-state rate of 2 percent per year, then we can calculate the present discounted value of the gain. This exercise reveals that the present discounted value of the growth effect is 0.47 percent of current income (or 0.46 percent if one uses the parameters from the export specification). By contrast the present discounted value of the level effect is 2.6 percent of per capita income: almost six times larger. Taken together the present discounted value of the gains from varieties that were newly imported between 1994 and 2003 is 3.1 percent of national income for the typical country.

Another way of thinking of these gains is the shift in permanent income. If we multiply these gains by the implied discount rate of 5 percent then we obtain the equivalent impact of these gains on permanent income. This exercise suggests that the gains from new imported varieties between the years of 1994 and 2003 raised permanent income of the typical country by 0.15 percent, of which, 0.13 percent is due to the level effect. The relatively small impact of the growth effect, even in the lab-equipment specification, suggests that most of the gains from new varieties arise from raising income levels.

VIII. Conclusion

How much does trade matter for growth? This is not a question that can be answered simply because trade regimes can vary enormously across countries and their impact may differ depending on a host of variables operative in the country carrying out reform. Rather than trying to answer a general, and perhaps not well-specified question, this paper has focused on quantifying one of the most important channels through which theorists believe trade affects growth. In particular, we estimate the impact that trade in new and better varieties has had on growth around the world. This is a central mechanism through which trade affects growth in many of the endogenous growth models and has never before been estimated. Moreover, rather than comparing aggregate measures of trade across goods or examining particular sectors in specific countries, we use a structure rich enough to allow for important differences across

sectors and countries, but flexible enough to allow for simple aggregation over sectors. This enables us to quantify the impact that new imported varieties has had on the global economy.

Our results indicate that while there is a lot of heterogeneity in the impact of new varieties on productivity growth, the typical estimated impact of new imported varieties on TFP is 10 percent as large as productivity growth in the typical country and a quarter as large as productivity growth in the typical developing country. We view these impacts as evidence that globalization is having an important impact on global growth rates.

Secondly, we provide evidence that the rise in world varieties is consistent with the predictions of the semi-endogenous knowledge driven growth model, and is not consistent with exogenous growth models or the comparative statics of continuum of goods or monopolistic competition trade models. This indicates that not only do new varieties have substantial impacts on world productivity, but R&D is associated with increases in exported varieties in the way assumed by the theory. Although most of the impact of new varieties on growth arises through the level effect, we estimate that for the typical country in the world, the combined effect of new imported varieties between 1994 and 2003 raised its permanent income by 0.2 percent. The present discounted value of these gains equals 3.1 percent of national income. These seem quite substantial.

Nevertheless, we need to mention several limitations of our approach. First, although our modeling of economic structure in any individual market is vastly richer than what finds in the typical macroeconomic analysis, it falls short of the careful empirical studies that can be produced using micro data. For any of the 200,000 markets that we consider in this paper, one could easily imagine more careful analyses of exactly how varieties affect welfare than the simple market structure we impose. Our decision not to take more care to model each of these sectors stems involves a trade-off of feasibility versus efficiency. No doubt more can be said about the export supply of Gambian raw, groundnuts than what we have produced in this paper, but we feel that the time necessary to do the careful analyses of markets and produce an estimate of the impact of trade on growth for more than a handful of countries makes this approach impractical for now.

Second, our close attention to the theory is both a strength and a weakness of this paper. Endogenous growth models themselves employ highly stylized descriptions of consumption, production, and innovation that do not have firm foundations in microeconomic studies of

consumer behavior or innovation. Leading macroeconomists have called for taking their models seriously and apply them to data, and we have tried to do so. However, one can have legitimate concerns about the underlying assumptions of the models themselves. In the future one can imagine developing richer models that better describe reality and better empirical methods that can take these models to the data. Nevertheless, we feel that much is learned by taking state of the art models to the data, and seeing how well they describe reality. We hope that is what we have accomplished here.

References

- Armington, Paul, "A Theory of Demand for Products Distinguished by Place of Production", International Monetary Fund Staff Papers, XVI (1969), 159-78.
- Backus, David K., Patrick J. Kehoe, and Timothy J. Kehoe, "In Search of Scale Effects in Trade and Growth," Journal of Economic Theory, LVIII (1992), 377-409.
- Broda, Christian, David Weinstein, "Variety Growth and World Welfare", American Economic Review Papers and Proceedings, XCIV (2004), 139-145.
- Broda, Christian, David Weinstein, "Globalization and the Gains from Variety", Quarterly Journal of Economics, forthcoming.
- Dorfman, Robert, Paul Samuelson, Robert Solow, Linear Programming and Economic Analysis, (New York, McGraw-Hill, 1958)
- Dornbusch, Rudiger; Fischer, Stanley; Samuelson, Paul A, "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods," *American Economic Review*, vol. 67, no. 5, Dec. 1977, pp. 823-39
- Eicher, Theo S. & Turnovsky, Stephen J. (1999) "Non-scale Models of Economic Growth," *The Economic Journal* **109** (457), pp. 394-415.
- Ethier, W.J., "National and International Returns to Scale in the Modern Theory of International Trade", American Economic Review, LXXII (1982), 389-405.
- Feenstra, Robert, "New Product Varieties and the Measurement of International Prices", American Economic Review, LXXXIV (1994), 157-77.
- Feenstra, Robert, Hiau Looi Kee, "On the Measurement of Product Variety in Trade", American Economic Review, XCIV (2004), 145-149.
- Feenstra, Robert, Chi-Yuan Liang, Dorsati Madani, Tzu-Han Yang, "Testing Endogenous Growth in South Korea and Taiwan", Journal of Development Economics, LX (1999), 317-341.
- Feenstra, Robert, James Markusen, "Accounting for Growth with New Inputs", International Economic Review, XXXV (1994), 429-447.
- Funke, Michael, Ralf Ruhwedel, "Product Variety and Economic Growth: Empirical Evidence for the OECD Countries," International Monetary Fund Staff Papers, XLVIII (2001), 225-242.

- Frankel, Jeffrey, David Romer, “Does Trade Cause Growth?”, American Economic Review, LXXXIX (1999), 379-399.
- Grossman, Gene, Elhanan Helpman, “Product Development and International Trade”, Journal of Political Economy, XCVII (1989), 1261-83.
- Grossman, Gene, Elhanan Helpman, “Trade, Innovation, and Growth”, American Economic Review, LXXX (1990), 86-91.
- Hall, Robert and Charles I. Jones (1999). “Why Do Some Countries Produce So Much More Output per Worker Than Others?,” *Quarterly Journal of Economics* February 1999.
- Hallak, Juan C. and James Levinsohn, “Fooling Ourselves: Evaluating the Globalization and Growth Debate,” NBER Working Paper #10244, January (2004).
- Harrison, Ann, “Productivity, imperfect competition and trade reform : Theory and evidence”, Journal of International Economics, XXXVI (1994), 53-73.
- Hummels, David, Peter Klenow, “The Variety and Quality of a Nation’s Exports”, American Economic Review, XCV (2005), 704-723.
- Jones, Charles, “Time Series Tests of Endogenous Growth Models”, Quarterly Journal of Economics, CX (1995), 495-525.
- Kocherlakota, Narayana, Kei-Mu Yi, “A Simple Time Series Test of Endogenous vs. Exogenous Growth Models: An Application to the United States”, Review of Economics and Statistics, LXXVIII, (1996), 126-134.
- Krugman, Paul, “Scale Economies, Product Differentiation, and the Pattern of Trade”, American Economic Review, LXX (1980), 950-59.
- Levinsohn, James, “Testing the imports-as-market-discipline hypothesis,” Journal of International Economics, XXXV (1993), 1-22.
- Noguer, Marta and Marc Siscart, “Trade Raises Income: A Precise and Robust Result,” *Journal of International Economics*, Vol. 65, Issue 2, March 2005
- Noguer, Marta and Marc Siscart, “FDI Trade and Income: A Causal Analysis,” Queens School of Business mimeo, 2006.
- Rauch, James, “Networks Versus Markets in International Trade” Journal of International Economics, XLVIII (1999), 7-35.
- Rivera-Batiz, Luis, Paul Romer, “Economic Integration and Endogenous Growth”, The Quarterly Journal of Economics, CVI (1991), 531-555.

- Rodriguez, Francisco and Dani Rodrik “Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence.” *Macroeconomics Annual 2000*, eds. Ben Bernanke and Kenneth S. Rogoff, MIT Press for NBER, Cambridge, MA, 2001.
- Romer, Paul, “Growth Based on Increasing Returns Due to Specialization”, American Economic Review, LXXVII (1987), 56-62.
- Romer, Paul, “Capital, Labor, and Productivity”, Brookings Papers on Economic Activity, Special Issue (1990), 337-67.
- Sala-i-Martin, Xavier. “I Just Ran Two Million Regressions,” *American Economic Review*, Vol. 87, n.2, pp. 178-183, 1997.
- Sato, Kazuo, “The Ideal Log-Change Index Number.” Review of Economics and Statistics, LVIII (1976), 223-28.
- Vartia, Yrjo, “Ideal Log-Change Index Numbers”, Scandinavian Journal of Statistics, III (1976), 121-126.

Figure 1
Growth in the Number of Goods Imported

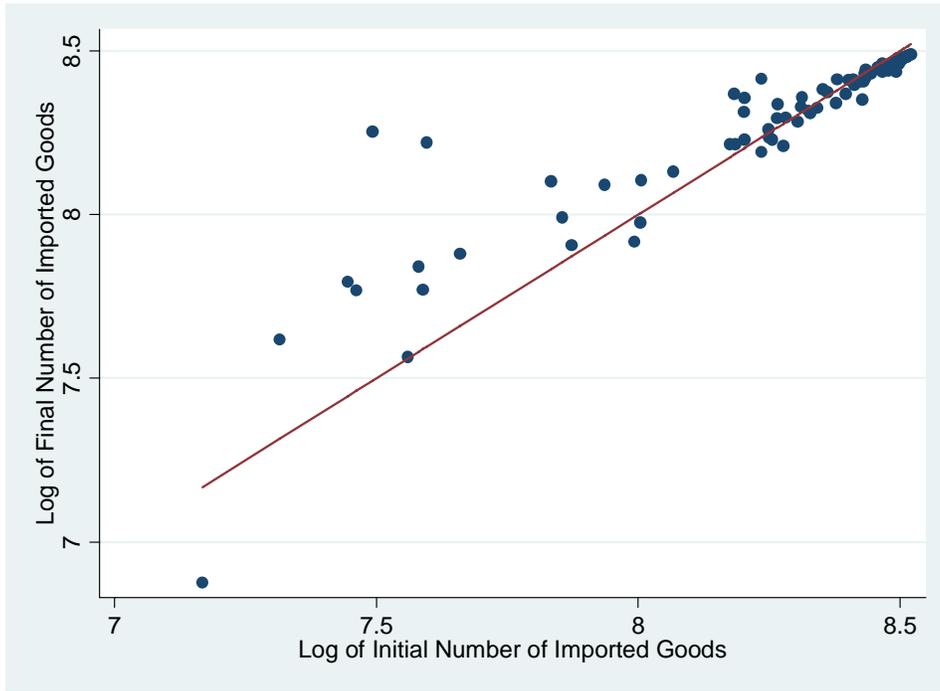


Figure 2
Growth in the Number of Source Countries per Good

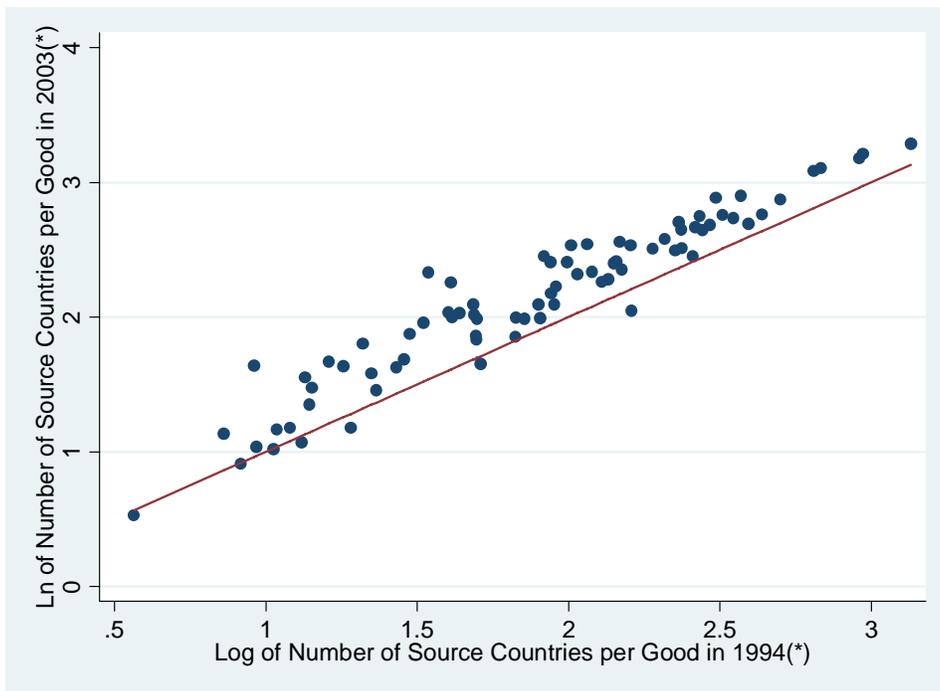


Figure 3

Growth in the Number of Imported Varieties

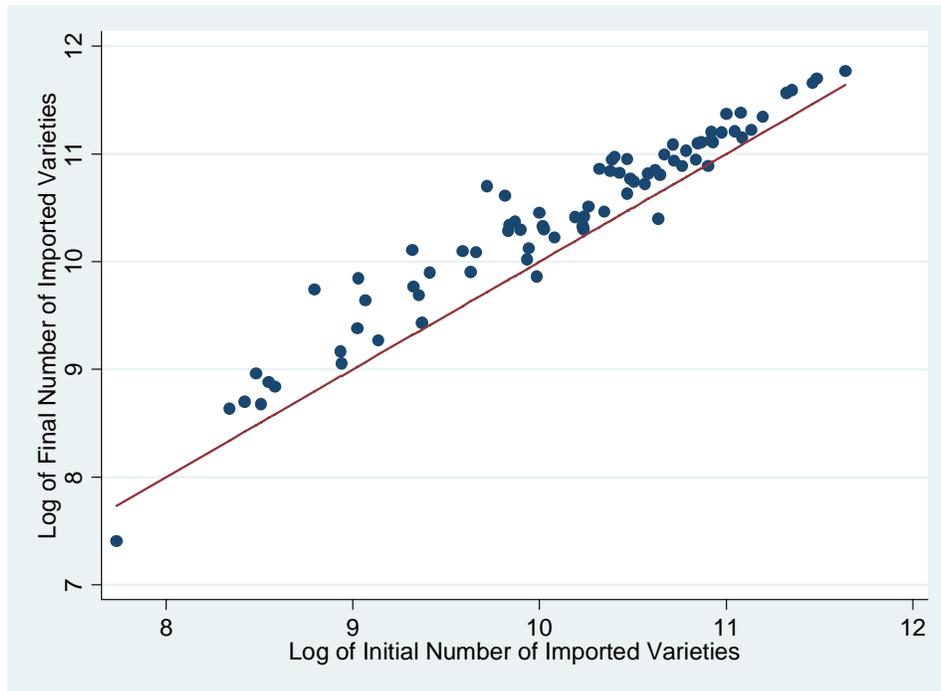


Figure 4

Frequency Distribution of the Number of Bilateral Trade Flows for Each Good

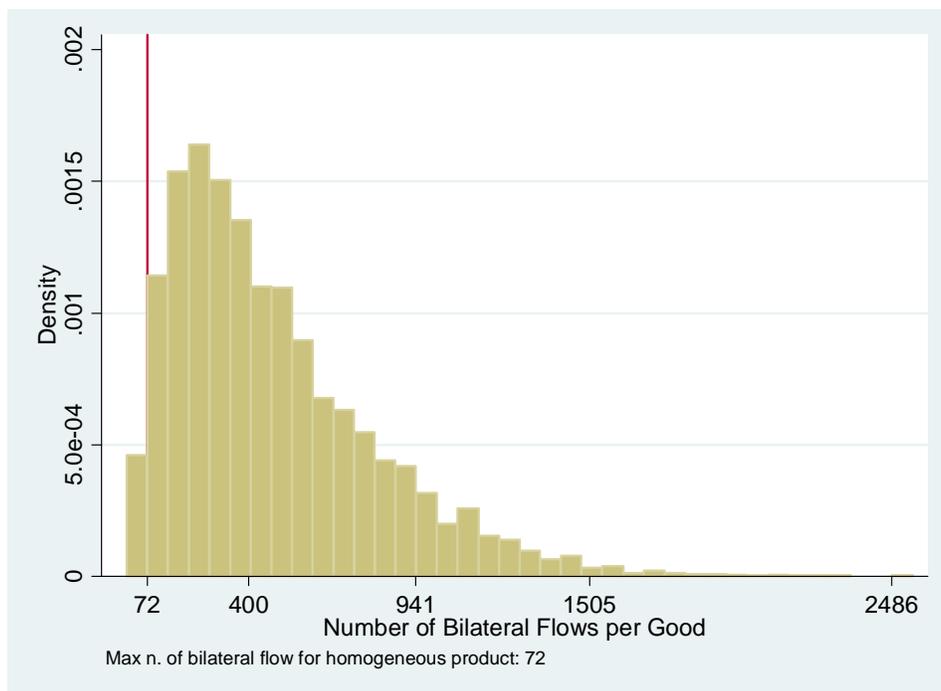


Figure 5

Number of Sectors with No Exports vs Size

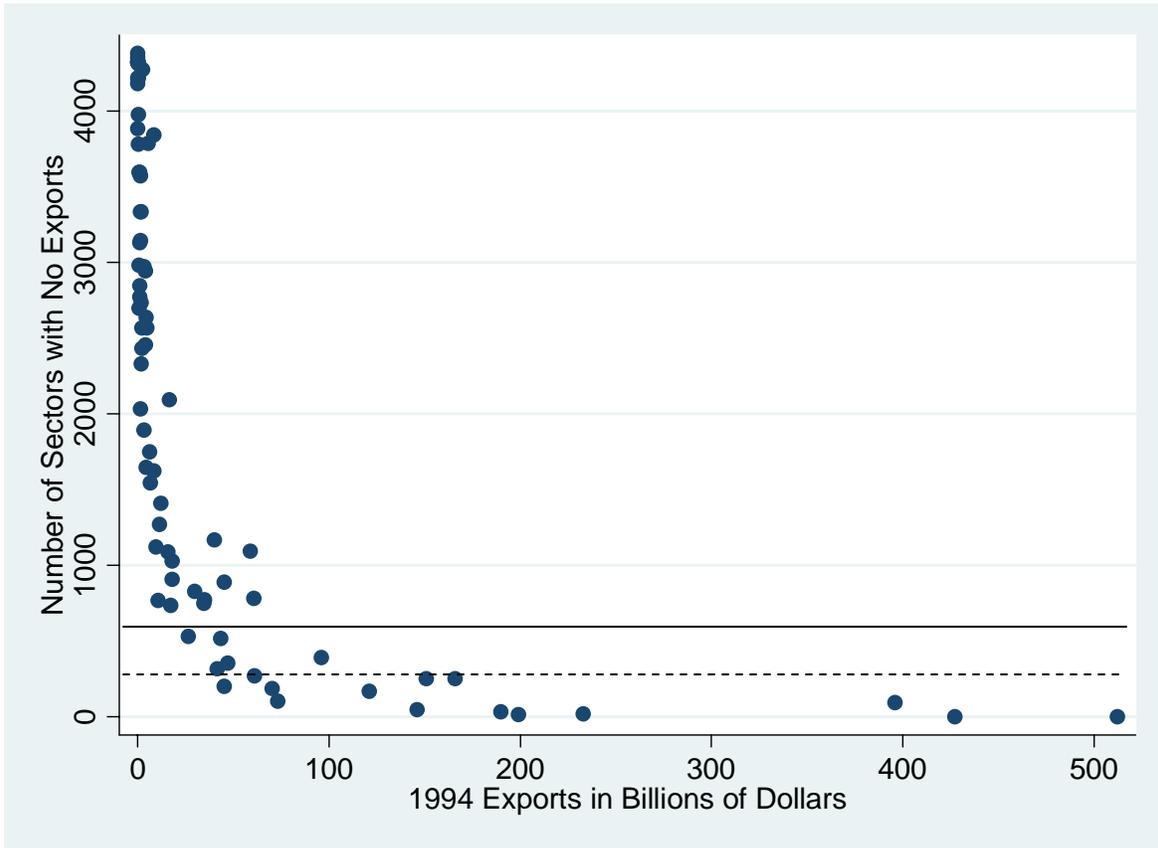


Figure 6

Output growth response after an exogenous 1.2 percent shock in Y
(Exogenous growth = 2%)

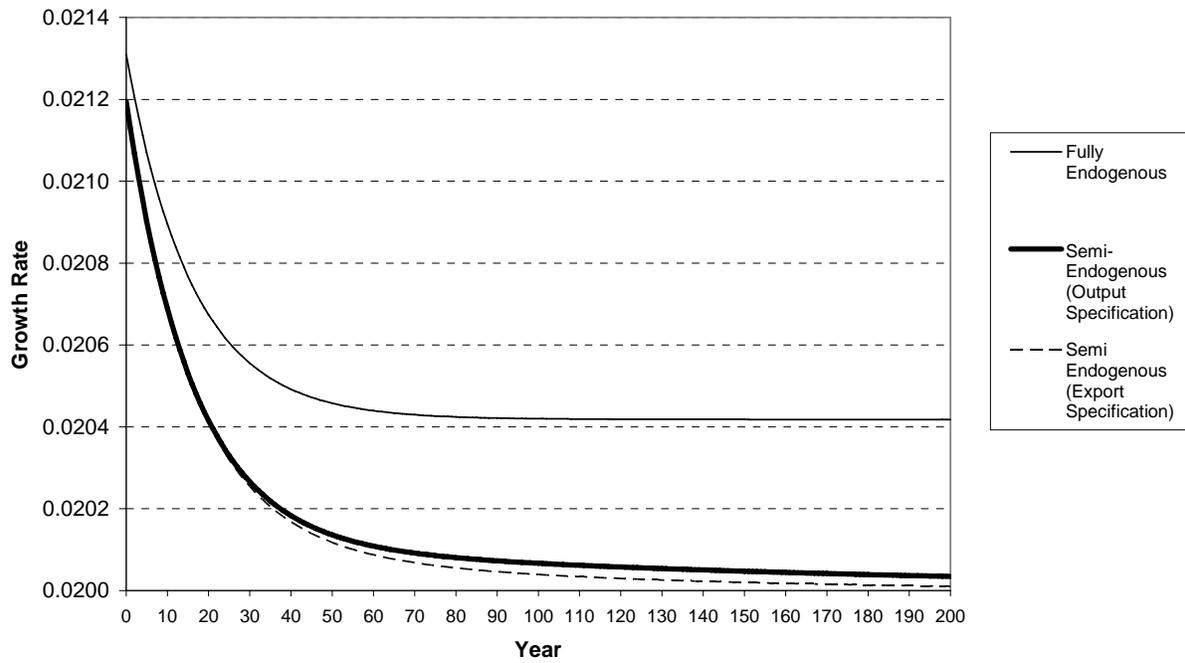


Table 1: Sample Statistics

	Median	Min	Max
First Year	1994	1994	1994
Number of Imported 6-digit Product Categories	4348	1297	5016
Average Number of Source Countries Per 6-digit	7.1	1.8	22.9
Number of Imported Varieties	32187	2281	113461
Final Year	2003	2003	2003
Number of Imported 6-digit Product Categories	4334	968	4960
Average Number of Source Countries Per 6-digit	9.8	1.7	26.8
Number of Imported Varieties	44484	1638	129420

Note: The 73 sampled countries are listed in the appendix table.

Table 2: Decomposing World Trade Growth in the 1994-2003

	Developed Countries	Developing Countries
Number of Countries	22	51
Median Growth in Trade/GDP	29.0%	16.2%
Median Growth in TradeCommon/GDP	15.0%	-16.4%
Median Growth in Imports	75.1%	62.5%
Median Growth inCommon Varieties	58.2%	36.9%
Median Growth inNew Varieties	21.4%	47.7%
Median Growth inDisappearing Varieties	-6.0%	-18.7%
Median Yearly Decay Rate of Common Varieties as a Share of Total Imports	0.7%	2.3%
10-year Decay Rate	6.6%	20.6%
20-year Decay Rate	12.8%	46.9%

Developed Countries are Australia, Austria, Canada, Denmark, Finland, France, Germany, Iceland, Ireland, Italy, Japan, Rep. of Korea, Mexico, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom, and the USA

Table 3: World Growth in Varieties

	1994	2003	Growth Rate
Total Imported Varieties	2,264,396	2,739,784	21%
Total Available Varieties	14,964,192	15,418,152	3%
Index of Trade Dispersion	0.140	0.167	19%

Table 4: Impact on Exact Import Price Index for 73 Countries 1994-2003

Rank in GDP per capita	Country Name	Median Sigma	St Error (Median Sigma)	Median Lambda Ratio	Per-year Productivity Gain (in percent)
1	Switzerland	3.8	0.15	0.965	0.053
2	Japan	2.8	0.18	0.977	0.024
3	Denmark	3.5	0.17	0.944	0.151
4	Norway	3.0	0.13	0.959	0.059
5	Germany	3.9	0.22	0.951	0.054
6	Austria	4.0	0.19	0.985	0.002
7	USA	2.3	0.13	0.984	0.024
8	Sweden	5.0	0.37	0.963	-0.047
9	France	3.7	0.21	0.904	0.094
10	Netherlands	3.3	0.19	0.877	0.534
11	Iceland	3.0	0.14	0.928	0.207
12	Finland	3.1	0.15	0.966	0.016
13	Hong Kong	4.2	0.24	0.991	0.131
14	Australia	2.5	0.10	0.978	0.133
15	Canada	5.0	0.42	0.977	0.057
16	United Kingdom	2.4	0.09	0.946	0.199
17	Italy	3.7	0.21	0.933	0.069
18	Ireland	3.8	0.16	0.946	0.228
19	Macau	4.3	0.44	0.984	0.076
20	New Zealand	3.2	0.21	0.963	0.130
21	Spain	2.8	0.14	0.945	0.136
22	Rep. of Korea	3.0	0.13	0.969	0.097
23	Cyprus	2.8	0.06	0.936	0.377
24	Greece	2.6	0.11	0.920	0.059
25	Portugal	3.4	0.15	0.956	0.066
26	Slovenia	3.7	0.21	0.924	-0.063
27	Saudi Arabia	2.9	0.11	0.942	0.094
28	Argentina	3.4	0.16	1.062	-0.012
29	Uruguay	3.4	0.19	1.006	0.244
30	Oman	3.9	0.18	0.504	1.163
31	Saint Kitts and Nevis	3.0	0.19	0.879	0.864
32	Gabon	3.3	0.18	0.930	0.958
33	Brazil	2.9	0.16	0.948	-0.007
34	Hungary	4.6	0.28	0.907	0.072
35	Chile	2.9	0.11	0.958	0.133
36	Malaysia	2.5	0.14	0.978	0.123
37	Croatia	5.0	0.37	0.886	0.286
38	Venezuela	3.3	0.17	0.916	0.118
39	Slovakia	4.0	0.23	0.903	0.374
40	Mexico	3.1	0.22	0.985	0.030
41	Mauritius	2.8	0.13	0.930	0.453
42	Belize	3.8	0.23	0.866	0.898
43	Poland	4.3	0.24	0.862	0.055
44	Dominica	3.1	0.12	0.877	0.897
45	Grenada	2.8	0.16	0.872	0.930
46	Thailand	2.9	0.20	0.945	0.167
47	Turkey	3.4	0.19	0.898	0.152
48	Saint Vincent	3.3	0.24	0.889	1.372
49	Colombia	2.9	0.18	0.963	-0.196
50	TFYR of Macedonia	3.3	0.21	0.800	0.618
51	Peru	3.1	0.16	0.943	0.151
52	Tunisia	3.2	0.16	0.936	0.339
53	Lithuania	3.8	0.23	0.627	0.582
54	Latvia	3.4	0.20	0.738	0.256
55	Ecuador	3.5	0.21	0.843	0.184
56	El Salvador	3.3	0.17	0.850	0.500
57	Jordan	3.1	0.14	0.925	0.441
58	Algeria	3.6	0.14	0.828	0.176
59	Romania	5.5	0.56	0.887	-0.061
60	Guatemala	3.3	0.19	0.907	0.072
61	Morocco	3.4	0.14	0.890	0.175
62	Egypt	3.7	0.22	0.975	-0.036
63	Indonesia	2.9	0.12	0.939	0.074
64	Bolivia	2.8	0.16	0.914	0.135
65	Sri Lanka	4.6	0.42	1.031	-0.216
66	Honduras	4.8	0.35	0.868	0.127
67	China	3.4	0.15	0.969	0.028
68	Nicaragua	3.9	0.37	0.877	0.889
69	India	3.7	0.21	0.858	0.097
70	Central African Rep.	3.7	0.49	0.990	0.367
71	Togo	3.1	0.20	0.826	0.350
72	Madagascar	3.6	0.27	0.722	0.869
73	Malawi	3.7	0.23	0.680	1.333
	Median Country World	3.4	0.19	0.93	0.13
	Median Developed	3.3	0.17	0.96	0.08
	Median Developing	3.3	0.19	0.91	0.18

Note: Developed Countries are defined in the notes to table 2.

TABLE 5
Table 5: Estimated Sigmas and Rauch Classification

	Rauch's classification of goods:		
	Commodity	Reference Priced	Differentiated
	<u>World 1994-2003 (3-digit goods, 6 digit varieties)</u>		
Mean across countries	12.1	7.3	7.2
Number of Observations		4139	7998
Test if different than Commodity (p-value)		0.000	0.000
Median across countries	3.8	3.0	3.3
Number of Observations		4139	7998
Test if different than Commodity (p-value)		0.000	0.000

Table 6: World TFP Accounting

	Per-year Productivity Gains due to Expanded Imported Varieties (1)	Per Capita GDP Growth Rates (2)	(1) as a share of (2) (3)	Per-year Productivity Gains due to Expanded Imported Varieties (4)	TFP Growth Rates (5)	(4) as a share of (5) (6)
	All Countries in Sample (N = 73)			All Countries in Sample (N = 53)		
Average	0.27 (0.04)	1.8 (0.5)	15	0.20 (0.03)	1.3 (0.2)	15
Median	0.13 (0.02)	1.7 (0.2)	8	0.13 (0.02)	1.1 (0.3)	12
	Developed Countries (N = 22)			Developed Countries (N = 19)		
Average	0.12 (0.01)	2.1 (0.1)	6	0.11 (0.01)	2.0 (0.2)	5
Median	0.08 (0.02)	1.9 (0.2)	4	0.09 (0.02)	2.0 (0.3)	5
	Developing Countries (N = 51)			Developing Countries (N = 34)		
Average	0.39 (0.05)	1.7 (0.6)	23	0.25 (0.05)	0.8 (0.2)	32
Median	0.17 (0.06)	1.4 (0.4)	12	0.13 (0.04)	0.6 (0.3)	22

Notes: All variables are expressed in percentages. Developed Countries are defined in the notes to table 2.

Table 7: Determinants of Innovation in New Varieties

Dependent Variable: Count of New Exported Goods

Measure of Size:	Exports	Output
<i>F</i>	0.27 [0.03]	0.27 [0.03]
<i>A</i>	0.38 [0.10]	0.69 [0.09]
N. Obs.	53	53
Pseudo R ²	0.09	0.09

Standard errors in brackets

Table 8: Growth and Level Effects

Impact on Permanent Income of Level Effect	Permanent Income effect of Temporary Rise on Output Growth Rate (after 9 years of variety growth)	Total Variety Impact = (1) + (2)
(1)	(2)	(3)
0.13	Typical Country in the World 0.02	0.15
0.08	Typical Developed Country in the World 0.01	0.09
0.18	Typical Developing Country in the World 0.03	0.21

Note: The difference in the growth effect in the output and export specifications only matters in the third decimal place.