

Dixit-Stiglitz approaches to international trade:†
Enough is enough

John P. Conley*
and
Robert Driskill**

May 2010

Revised: April 2013

† We thank seminar participants at Washington University, St. Louis, Academia Sinica, Taipei, Universidade Catlica Portuguesa, Lisbon, the ES meetings in St. Louis, the SAET Meetings in Faro, Portugal, the PED meetings in Bangkok, the 2011 Southern Economic Association meetings, the XIII Conference on International Economics at Granada, Spain, Eric Fisher, Benedetto Molinari and Manfred Dix for comments. We especially wish to thank Hideo Konishi for many long discussions and debates. We take full responsibility for any errors.

* Department of Economics, Vanderbilt University Nashville, TN 37235,
j.p.Conley@Vanderbilt.edu

** Department of Economics, Vanderbilt University Nashville, TN 37235,
Robert.a.driskill@vanderbilt.edu

Abstract

We show that the Melitz (2003) continuum interpretation of the Dixit-Stiglitz monopolistic competition model does not represent a reasonable limit of a large finite economy unless it is built on a micro-foundation of heterogeneous agents. We argue for both theoretical and empirical reasons that each of these heterogeneous agents must choose to consume at most a strictly bounded number of commodities instead of a strictly positive quantity of each of the infinite number of goods produced (as in Melitz). We show in a highly general framework that this implies that (1) the benefits to product diversity converge to zero as the economy gets large and (2) the proportion of imports goes to zero as countries get large in the presence of non-zero iceberg costs. We offer a new theoretical approach to product innovation that highlights embedded assumptions in Dixit-Stiglitz that we argue drive counter-factual results.

1. Introduction

The Dixit-Stiglitz (1977) model of Chamberlinian monopolistic competition model (hereafter, DS) has had an enormous impact on research in Industrial Organization, Economic Geography, Monetary and Real Business Cycle Theory, Growth Theory, and especially in International Trade. Beginning with Krugman (1979, 1980) and more recently extending through the work of Melitz (2003) and his followers, DS has become the workhorse model in the field.¹

Traditional theory predicted that countries with differences in technology (the Ricardian model) or with differences in endowments (the Heckscher-Ohlin model) were likely to find gains from trade. This suggested a pattern of trade in which the goods that a country imported would quite differ from those that it exported. Empirical studies by Balassa (1967), Grubel and Lloyd (1975), and others, however, showed that there was also significant trade among countries with similar technologies and endowments. Even more surprising, it was often the case that very similar items were being exported and imported. For example, Italy might export Fiats to Germany and import Volkswagens in return.² It was therefore attractive that the theoretical predictions that the DS approach to trade yielded seemed to be consistent with these empirical observations. Of course, it was also attractive that these DS models have closed form solutions and offer a flexible and mathematically tractable foundation for the theoretical study of how government interventions, changes in the underlying economic environment and other factors might be expected to affect the pattern of trade.

Unfortunately, this tractability came at the cost of certain simplifying assumptions. Most important of these was that the DS model assumed a representative consumer who benefits directly from product variety and who always chose to consume a strictly

¹ Melitz (2003) has over 3000 citations according to Google Scholar. Peter Neary (2004, p. 161) recently wrote: “It is even rumored that there are universities where the graduate trade curriculum covers nothing but monopolistic competition!”, and in a similar vein Broda and Weinstein (2006) write: “It is striking that in the quarter-century since Krugman (1979) *revolutionized* international trade theory by modeling how countries could gain from trade through the import of new varieties ...” (italics ours).

² See Neary (2004) for a more complete analysis.

positive level of all goods produced in the economy.³ In addition, DS approaches assumed that the economy was “large” in the sense that there were enough firms (each producing a distinct product) that producers would be price-takers and would therefore not act strategically.

The central point of this paper is to argue the first assumption is “unrealistic”, not just in the conventional sense that all economic models abstract from the real world, but in the sense that it fundamentally drives the theoretical conclusions and leads to predictions that are, in fact, inconsistent with empirical observations. If one reconstructs the model on what we argue is a more reasonable theoretical foundation, one reaches very different and theoretical conclusions which in turn have very different policy and empirical implications. In particular, one is forced to conclude that trade is likely to take place only in markets in which firms do have market power and so will not be price takers. In other words, the price taking assumption which is key to obtaining for a closed form solution in any DS model is fundamentally inconsistent with derived market equilibrium behavior for a large (or indeed, any) economy. In addition, if the only motivation for trade is a desire for product diversity, then the benefits from trade converge to zero as economies get large, and in fact, trade ultimately disappears completely in the presence of strictly positive iceberg costs.

The problem is most starkly displayed in the continuum version of DS-trade model as developed in Melitz (2003).⁴ As is well known, it is quite possible to write a mathematically correct measure theoretic model that suggests underlying structures that may

³ Preferences are generally assumed to satisfy Inada conditions and so all consumers demand a strictly positive quantity of every good regardless of either the number of goods or the price of any particular good. Although Melitz and Ottaviano (2008) consider a model in which not all varieties of differentiated products are consumed, the model still assumes all consumers are identical and thus does not address our central criticism. See Broda and Weinstein (2006) for the empirical implementation of the DS model in an investigation of the gains from trade due to increased variety, and again, see Neary (2004, p. 174) for a more complete description of the non-diminishing returns to variety feature.

⁴ It is interesting to note that in an early unpublished version, Dixit and Stiglitz (1974) did assume a continuum, only to abandon the approach in their next version, Dixit and Stiglitz (1975). In the 1975 version, they claimed in a footnote that they had abandoned the continuum version because “technical difficulties ... led to unnecessary confusion.” Recent treatments, however, almost exclusively use the continuum version, without apology. We will argue that although there are technical problems with the continuum approach, in fact, any model in which both large (but finite) numbers of goods are produced and agents are assumed to consume positive levels of all goods will give misleading conclusions, at least in the context of trade.

or may not be economically sensible. Unless the continuum economy one constructs is a reasonable limiting case of a large finite economy, one risks finding results that are only artifacts of the mathematical approach and thus do not give correct economic insights.⁵

With this in mind, consider Melitz’s model of international trade. His approach is to use a representative consumer whose preferences are given by a C.E.S. utility function over a continuum of goods indexed by ω :

$$U = \left[\int_{\omega \in \Omega} q(\omega)^\alpha d\omega \right]^{1/\alpha}$$

where the measure of the set Ω represents the mass of available goods. Then $P = \left[\int_{\omega \in \Omega} p(\omega)^{1-\sigma} d\omega \right]^{1/1-\sigma}$ is the price of the “aggregate good” Q where $\sigma = 1/(1-\alpha) > 1$ and:

$$q(\omega) = \left[Q \frac{p(\omega)^{-\sigma}}{P} \right].$$

Note that this means that strictly positive levels of each good are consumed by the representative agent. The question, therefore, is to understand what sort of large finite economy the Melitz model represents as a limiting case. We see two possibilities.

Perhaps the economy consists of an infinity of homogeneous agents with utility functions exactly like the representative agent above. The problem with this is quite simply that it is resource infeasible. To see this, suppose that the good with index number .3 is a banana and each agent consumes one banana. If we integrate to find the number of bananas consumed in total, a measure one of bananas must be produced. If it costs one unit of labor per banana (for example), a measure one of labor is required. Thus, the measure of bananas consumed is of the same order of magnitude as the total

⁵ A classic example of how the mathematical features of the continuum can yield artificial results can be seen in the marriage problem. Suppose there is an interval $[0, 1]$ of girls, and $[0, 2]$ of boys. If we match girl number $g \in [0, 1]$ to boy number $2g \in [0, 2]$, each and every agent has exactly one partner. Clearly, this is nonsense. If the large finite economy has twice as many boys as girls, then half the boys must remain unmatched. The continuum model, in other words, does not reflect the economic fact that girls are scarce. There is little predictive value in the continuum model in this case. (See Kaneko and Wooders 1986 on “Measure Compatibility” for a solution.)

supply of labor. Unfortunately, since all agents consume a strictly finite amount of each of an (uncountably) infinite number of goods, an infinite measure of labor is required. Needless to say, this is a rather serious violation of the resource constraint. Even if this were feasible to produce, agents don't have the physical ability to consume a finite amount of an infinite number of goods.⁶

Note that it is also true that in an economy with a large but finite number of both consumers and goods, per capita consumption could never be strictly bounded from zero as both the number of goods and agents increased for exactly the same reasons. The core of the problem can really be traced back to Dixit and Stiglitz who embedded in the structure of their model that all agents consume some of all goods produced in the economy. Of course the Dixit and Stiglitz approach does not violate the budget constraint since the consumption levels of each good go to zero as the number of goods increase. If one takes this to the continuum limit as Melitz does, one would therefore expect to see all agents consuming infinitesimal quantities of an infinite set of goods. However, this does not appear to be what Melitz has in mind. To state such a model correctly, we would have to resort to a non-standard analytic approach to reflect that fact that demand for each good would literally be infinitesimal, but when multiplied by prices and integrated, expenditure would still equal the representative agent's finite income. Even if one went to the trouble to follow this approach, it is doubtful that the correct (meaning economically relevant) limiting case of an economy with a large number of goods is for agents to choose to consume smaller and smaller slivers of all the goods the economy produces. From this we conclude that it is hard to give a reasonable interpretation of the Melitz model if one insists on both homogeneous agents and an infinite or even large number of goods.

Alternatively, perhaps the utility function is supposed to be a reflection of aggregate behavior and so the resulting demand is really meant to be interpreted as market

⁶ Another way to see this is that ask: what would it cost an individual to buy a finite quantity of an infinity of goods? If the prices (which relate to input costs) are also finite, then the cost is infinity which is, of course, a slight violation of agents' budget constraints. We wish to emphasize that we are in no way suggesting that there is a mathematical flaw in the Melitz model. We are only pointing out that it is not a reasonable limiting case of a large finite economy.

demand. This would solve the feasibility issue, but in this case, the “aggregate” utility function Melitz considers *cannot* be representative of a homogeneous and infinite set of consumers (or by extension, a homogeneous set of consumers in economy with a large but finite number of agents and goods). If consumers were homogeneous, then each of them would have to consume some of each good produced and we would be back to the first case discussed above. Thus, to be sensible, the utility function must represent the aggregate behavior of an infinite set of *heterogeneous* consumers. This means there must be some micro-foundation with heterogeneous agents behind the Melitz approach. The correct interpretation of Melitz’s model and results therefore depends upon exactly what this micro-foundation is.

The main point of this paper is to explore what properties such a micro-foundation needs to have and to discover the implications for policy. We begin with a key observation: In real life, individual agents (even in a large finite economy) consume at most a finite (in fact, a strictly bounded) number of goods. For example, how many of the goods offered at Amazon have you personally consumed? One reason that you only purchase a tiny fraction of Amazon’s offerings might be transactions costs. Even if you had infinite money and stomach capacity, you would not have the time to complete an infinity of transactions. A second reason might be non-convexities driven sometimes by indivisibilities. Half a Camry plus half an Accord will not take one as far or as fast as a whole Maxima.⁷

We therefore add explicit constraints that imply that each agent will choose to consume only a strictly bounded number of goods. This means both that agents must choose a subset of all the goods produced in the world to consume and that firms must anticipate how many agents will choose to add their products to their consumption sets in addition to guessing the demanded quantities. This in turn leads to important

⁷ We find it interesting that one of the most salient examples used to illustrate the idea of intraindustry trade is in fact that of the auto industry, an example for which the idea of consuming a “little bit of everything” is obviously inappropriate. For example, in the best-selling undergraduate textbook of Krugman and Obstfeld (2009), the “Case Study” for “Intraindustry trade in action” is the North American Auto Pact of 1994, (p.134-5), and the numerical problem they give concerns a fictitious auto industry (problem 5, p. 150).

questions about the nature of innovation and product differentiation. In the DS framework, all products are identically desirable. In this sense, each potential new product a firm might choose to produce is perfectly substitutable for each existing product. In contrast, we argue that product innovation is a significant force that drives both trade and consumer welfare. What fraction of goods that we buy today even existed twenty years ago, for example? We therefore develop two different modeling approaches that allow products to be truly distinct and which allow us to think about both consumption and product development choices in a non-trivial way.

Our work, then, is part of a long-standing literature that criticizes the DS framework as one with important limitations. We think our major contribution is to show that these limitations are especially important in providing misleading lessons about the gains from trade and the pattern of trade in relatively large economies. The rest of the paper proceeds as follows. In section 2, we follow Melitz and formally show how heterogeneous consumers are required for the model to make sense. In section 3 we propose a more general model and prove some theorems. In section 4, we make some remarks the more general implications of approach suggest for product innovation. Section 5 concludes.

2. An Algebraic Example with Bounded Consumption

In this section, we will dispense with the representative agent approach and instead build a model based on heterogeneous agents each of whom is limited in the number of goods that they can simultaneously consume.

Before doing so, we should note that we are not the first to suggest that the DS “representative consumer” approach has problems. For example, Perloff and Salop (1985) point out that the DS approach would not apply to the many types of products for which we usually observe consumers purchasing only one (or a few) unit(s), such as automobiles. Hart (1979, 1985a, 1985b), on the other hand, emphasized problems with

the limiting behavior of the DS model. All of these papers argued that it was necessary to incorporate heterogeneous consumers into a monopolistic competition model. Their focus was on the foundations of monopolistic competition, however, and not on the extensions of the DS model to trade and the implications of those extensions for welfare gains from product variety.

In the trade literature itself, Lancaster (1980) introduced the “ideal” variety approach as a method for explaining intraindustry trade among similar countries. The tractability of the DS model seems to have limited the appeal of this approach to modeling trade and variety.⁸

More recently, Mrazova and Neary (2011) note that almost all results on the selection of more productive firms into export activities have been developed within the DS monopolistic competition paradigm. This basic idea suggests that more efficient firms should also be large, which is not consistent with the idea of monopolistic competition. Thus, Mrazova and Neary argue there is an essential tension between the assumption of price taking and the result that efficient firms tend to grow larger in equilibrium. This idea finds strong support in the models we develop below.

At a formal level, assume that there are an uncountably infinite set of heterogeneous consumers in the economy and denote agents by the index $i \in [0, I] \equiv \mathcal{I}$. We assume the interval of consumers is endowed with Lebesgue measure. Agents are each endowed with an amount ω of a “traditional” good denoted y .

There is a countably infinite set of potential goods denoted produced by firms which arise endogenously in the economy. Each firm that decides to enter the market uses the traditional good to produce one and only one commodity. We denote the production of a typical firm as X_n for $n \in \mathcal{N}$ where \mathcal{N} is the set of goods produced in the economy.⁹ We will often use $N = |\mathcal{N}|$ to denote the number of goods that firms

⁸ As Neary (2004, p. 160) puts it: “Ultimately, though, these alternative approaches proved less tractable and hence less fruitful than the DS specification.”

⁹ In general, \mathcal{N} may not be the sequential index set $\{1, \dots, N\}$ since any set of goods might in principle be optimal for the firms to produce. In fact, \mathcal{N} need not be finite, in general, although it will be in equilibrium in the economy we describe below.

produce in equilibrium. On the cost side, firms must spend a fixed amount to enter production and then a variable amount that is proportional to the quantity produced. Thus, firms have identical affine cost functions of the following form:

$$C_n(X_n) = F + VX_n,$$

where $F, V > 0$ are technological parameters, and $C_n(X_n)$ denotes the measure of traditional good needed to produce a measure X_n of good n . This means that if x_{ni} is the level of good n consumed by agent i , then material balance requires that $X_n = \int_i x_{ni} di$. Note that we construct the model so that fixed costs are large compared to size of agents in the sense that it will take a positive measure of labor cover these costs. This is important because if fixed costs were something finite (that is, could be paid for by the labor of a finite set of workers), they would be negligible in proportion to the variable costs of production if the firms produced positive measures of goods. Moreover, this reflects the idea that fixed costs are somehow related to capacity. For example, a marketing and distribution network might be fixed costs that are of the same scale as size of the market one wishes to serve.

We take a simple approach to modeling the idea that agents have a limited capacity to consume goods in this section. While many goods may be offered, we assume that there is an exogenously given upper bound, K , on the number that can be consumed by an agent at any one time. Agents must therefore choose a set $\mathcal{N}_i \subset \mathcal{N}$ where $\|\mathcal{N}_i\| \leq K$. We call this the *demand set of agent i* . A typical good is denoted x_n . We assume the following functional form for agents' utility functions:

$$U_i(x, y) = \sum_{n \in \mathcal{N}_i} \frac{(\alpha_{ni} x_{ni})^{\frac{1}{2}}}{\frac{1}{2}} + y_i$$

The essential element of our approach in modeling this choice is that for each $i \in \mathcal{I}$ and each $n \in \mathcal{N}$, the preference parameter, α_{ni} , is independently drawn from a uniform distribution over $[0, 1]$. We assume that these heterogeneous agents are fully informed about the set of commodities being produced and their own preference

parameters associated with each. Note that in contrast to the DS, this approach to product innovation does not imply new products are equally desirable to all agents. As new products are developed, agents separately evaluate them and may find that they are desirable or undesirable substitutes for existing products.

Given this, the problem of the consumer is as follows. Normalizing the price of the endowment good to $p_y = 1$ and taking the demand set as given for the time being, we get the following budget constraint for agent i :

$$y_i = \omega - \sum_{n \in \mathcal{N}_i} p_n x_{ni}$$

Substituting this into the preferences gives us the following problem:

$$\max U_i(x_i, y_i) = \frac{(\alpha_{ni} x_{ni})^{\frac{1}{2}}}{\frac{1}{2}} + \omega - \sum_{n \in \mathcal{N}_i} p_n x_{ni}$$

This gives the following first order conditions:

$$\frac{\partial U}{\partial x_{ni}} = -p_n + (\alpha_{ni})^{\frac{1}{2}} (x_{ni})^{\frac{1}{2}-1} = 0, n \in \mathcal{N}_i$$

Solving this gives:

$$(x_{ni})^{\frac{1}{2}-1} = \frac{p_n}{(\alpha_{ni})^{\frac{1}{2}}} \left((x_{ni})^{\frac{1}{2}-1} \right)^{\frac{1}{2}-1} = x_{ni} = (\alpha_{ni})^{\frac{1}{2}} (p_n)^{\frac{1}{2}-1}$$

This in turn gives the following demand system for agent i :

$$x_{ni} = \frac{(\alpha_{ni})}{(p_n)^2}.$$

Up to this point, we have taken the demand set as given. To find the aggregate demand for any specific good n , however, we need to know how many agents choose to put this good in their demand sets. Suppose we assumed all firms behaved as price-takers and that in equilibrium, all firms set equal prices. Then agents would

choose to divide their income over the K goods for which their preference parameters happened to be the highest (since this would clearly yield the highest utility).

Notice that this creates an interesting dynamic as more firms enter the market. Each time a new product is created, each agent draws a new utility parameter between zero and one, and if it happens to exceed lowest parameter currently in his demand set, the agent replaces the lowest parameter good with the newly created good. Thus, as new goods continue to enter the market place, the agents are able to get more picky and only goods with very high alphas are retained in their demand sets. Eventually, all the alphas in each agent's demand set approach one, the highest possible value. Intuitively, one can see that if many goods are already being offered by firms, the benefit to agents of even more product diversity diminishes. Agents already are close to having ideal goods (with alphas approaching one) and getting ever so slightly better goods yields comparatively little additional benefit. This intuition will be made rigorous in the next section.

In any event, since we know that each agent has a utility parameter for each produced good that is independently drawn from a uniform distribution over the unit interval, it is possible to use a combinatoric argument to calculate exactly what the aggregate demand for each good would be and also how many firms would enter in equilibrium.¹⁰

Unfortunately, since there will be a finite (though possibly very large) number of firms producing goods in equilibrium, firms will not, in fact, be price-takers. If we took the realistic step of also assuming that costs are heterogeneous and randomly drawn, then equilibrium prices also would not be symmetric in general. Including these factors turns out to make getting a closed form solution difficult or impossible. We will therefore set aside the attempt to solve an algebraic model for the time being, and consider the problem from a general standpoint.

In the next section we will take up a model that is far more general than the one

¹⁰ Contact the authors for an earlier draft of this paper in which this exercise is carried out.

explored by Melitz. Our central findings is that in general, only a finite number of goods would ever be produced regardless of the size of the economy. In essence, we show that even very large economies become satiated in product variety just as individuals become satiated in individual goods. Our results therefore suggest that trade is most likely to take place when there are relatively few different product types being offered in a sector (think aircraft and automobiles). Firms in these sectors clearly behave as oligopolists, not as price-taking competitors. We argue that using a model that has price-taking firms at its foundation as a basis for suggesting trade policy is therefore dangerous. One would not try to regulate the airline, telecommunications or pharmaceutical industries as if these sectors were competitive. Similarly, if significant fractions of cross-hauled goods are produced by firms with few competitors, then policy and regulatory conclusions for trade that ignore this fact are likely to be quite misguided.

3. General Model

In this section, we consider the planner's problem for a much more general economy. We retain the quasi-linearity of agents' utility, but do not impose a functional form or even convexity or monotonicity on the utility functions. We drop the *ad hoc* limit on the number of goods that agents can add to their demand sets and instead we make this an endogenous choice on the part of agents. Agents could in principle choose to consume all the goods produced, even an infinite number. Quasi-linearity is convenient here since it provides us with a clear metric of social welfare. If we were willing to state a welfare function, on the other hand, we could dispense with this assumption and the results given below would continue to hold, even if we added additional non-produced goods.

On the production side, we retain the affine production technology, but assume that firms are heterogeneous and randomly realize different fixed and variable cost param-

eters after paying a research and development cost. Firms may decide not to produce after seeing their *ex post* cost functions. Again, it would be possible to generalize these results to have several non-produced goods as inputs with arbitrary technology. All that would be required for the results to continue to hold is that there be a lower bound on the implicit fixed and marginal costs. We don't feel much insight is added from this generalization so we consider the simpler case.

Our strategy here is not to solve the exact planner's problem. This turns out to be quite complicated since we would have to maximize the expected payoff from choosing any number of development projects, which in turn would require calculating the maximum payoff from choosing any subset of these developed projects for any realization of fixed and variable costs, and then choosing the maximizing production levels and allocations of goods across agents. Given the lack of convexity and the minimal assumptions on utility, this is quite difficult and it might not be the case that the implicit optimal allocations would even be measurable. Instead, our strategy is to construct a series of well defined and feasible plans. Since the plans are feasible, each one provides a lower bound on the social welfare one would obtain in any socially optimal plan. We then show that as the economy gets large, the welfare obtained in these "guarantee point" plans converges to maximum possible social welfare one could ever obtain. Since we show that there is a uniform finite upper bound on the number goods produced in each of these guarantee point plans, we conclude that the social payoff to increasing the number of goods produced diminishes to zero in the limit. Finally, we show that this in turn implies that volume of trade between any two economies diminishes to zero as both of the economies get large, in contrast to the finding of DS based approaches.

More formally, we consider an economy with one non-produced endowment good y which is initially held by agents. All the other goods are produced by firms using the endowment good. We take the following approach to product innovation which is new as far as we are aware, although it is related to the "public project" approach to public goods equilibrium pioneered by Mas-Colell (1980). Produced goods each have a

product characteristic g drawn from separable metric space (G, d) where G is a convex and compact set.¹¹ We will call G the *innovation space*. We will discuss implications of this “innovations space” approach at more length in section 4. Since we focus on the social planner’s problem in this section, we abstract from the issue of how firms decide to enter and exit. The planner, therefore, decides how many development projects to undertake. To make the argument simpler, we will assume that the planner can choose exactly which goods to produce but that research and development costs are D per project. We denote the set of goods produced by the planner as \mathcal{N} .

The cost of producing a good g is given by an affine function:

$$C^g(Y^g) = F^g + V^g Y^g,$$

where F^g and V^g are independently drawn the intervals $[\underline{F}, \overline{F}]$ and $[\underline{V}, \overline{V}]$ under frequency distributions f^F and f^V , respectively where $\underline{F}, \underline{V} > 0$.

We assume that there is there an interval of agents $[0, I] \equiv \mathcal{I}$ endowed with Lebesgue measure. Agents are each endowed a strictly bounded quantity of the non-produced good: $\omega_i \leq \overline{\omega} > 0$. Let $\Omega(i)$ be a measurable function that gives the endowment for the economy.

We assume that agents can always consume the non-produced good, but deciding to consume a positive quantity of any produced good, and thus, adding it to their utility functions, is costly. This might be due to a transactions cost of ordering a quantity of a new good or of going to a store to buy some, or perhaps it might reflect an attention or effort cost of deciding to start to consume a new good.¹² Since it will not matter for the point we are investigating in this section, we will take the simple approach that if an agent i decides to add a produced good g to his demand set, \mathcal{N}_i , he must pay a cost of $t > 0$ of the non-produced good to cover the transactions costs. As a result, the size

¹¹ The space of potential product characteristics might be a finite dimensional Euclidean space, but need not be. It is important to note that the meaning of this is if two goods are close under d they are *physically* similar. That they are close utility substitutes is therefore an implication of this fact, not an separate assumption.

¹² More fundamentally, this might be due to a time cost of engaging in the act of consuming a good.

of the demand set $|\mathcal{N}_i| = N_i$ is endogenously determined and may differ across agents. Given this, the utility function of the agent i is given by the quasi-linear function:

$$U_i(x_i, y_i) = y_i + h_i(\{x_i^g\}_{g \in \mathcal{N}_i}) - tN_i$$

We make two key assumptions on agents and goods.

Assumption A. Agents in a close neighborhood of each other are similar: For

all $\epsilon_{\mathcal{I}} > 0$ there exists $\delta_{\mathcal{I}} > 0$ such that for all $i, j \in \mathcal{I}$ such that

$$|i - j| < \delta_{\mathcal{I}}$$

$$(a) \quad |\omega_i - \omega_j| < \epsilon_{\mathcal{I}}$$

(b) for any demand set $\mathcal{N}_i \subset G$ and consumption bundle (x_i, y_i) permissible under \mathcal{N}_i , $|U_i(x_i, y_i) - U_j(x_i, y_i)| \leq \epsilon_{\mathcal{I}}$

This says that if agents are close together as measured by their identity index, then they have similar endowments and get similar utility levels from the same consumption bundles.¹³

Assumption B. Goods in a close neighborhood of each other are close substitutes:

For all $\epsilon_G > 0$ there exists $\delta_G > 0$ such that for all $i \in \mathcal{I}$ and all integers N :

(a) for all $\overline{\mathcal{N}}_i = (\bar{g}_1, \dots, \bar{g}_N)$, $\widehat{\mathcal{N}}_i = (\hat{g}_1, \dots, \hat{g}_N)$ such that for $n = 1, \dots, N$, $d(\bar{g}_n, \hat{g}_n) < \delta_G$ and $x_i^{\bar{g}_n} = x_i^{\hat{g}_n}$ then $|U_i(\bar{x}_i, y_i) - U_i(\hat{x}_i, y_i)| < \epsilon_G$.

(b) for all $\overline{\mathcal{N}}_i = (\bar{g}_1, \dots, \bar{g}_N)$, if for some $\bar{g}_n \in \overline{\mathcal{N}}_i$, $\widehat{\mathcal{N}}_i = \overline{\mathcal{N}}_i \setminus \bar{g}_n$ and for some $\bar{g}_m \in \overline{\mathcal{N}}_i$ $d(\bar{g}_n, \bar{g}_m) < \delta_G$ it is the case that $x_i^{\hat{g}_m} = x_i^{\bar{g}_m} + x_i^{\bar{g}_n}$ while for all $k \neq n, m$, $x_i^{\bar{g}_k} = x_i^{\hat{g}_k}$ then $|U_i(\bar{x}_i, y_i) - U_i(\hat{x}_i, y_i)| < \epsilon_G$.

This says that if we take an arbitrary bundle of goods and replace each of these goods with an identical levels of very similar goods as measured by the characteristic space metric d , then utility levels change very little. In addition, if two goods in a

¹³ This is less general than allowing agents to have arbitrary utility functions, but is far more general than assuming that agents are identical or have utility functions with particular functional forms, as Melitz seems to.

consumption bundle are similar and we force an agent to consolidate this consumption on one of the goods rather than splitting it between the two, utility levels change very little.

Next, we make two minimal assumptions on the utility function:

Assumption C. Continuity and diminishing utility of produced goods: For all $i \in \mathcal{I}$, h_i is continuous, differentiable, and there exists an upper bound \bar{B} such that for any consumption vector (x, y) that includes some good g for which $x^g > \bar{B}$, it holds that: $\partial h(x)/\partial x^g < \underline{V}$

Note that convexity and monotonicity are not needed. We only require that the marginal willingness to pay for any produced good falls below the minimal marginal cost of production at some high consumption level. Without this, it might be possible for an agent to get infinite utility by allowing the consumption level of the endowment good to be infinitely negative in order to pay or infinitely high levels of produced good in which he is never satiated.

Finally, we assume there is at least a positive probability any given pair of cost parameters will be drawn for any particular g .

Assumption D. Full support in cost space: The frequency distributions f^F and f^V over $[\underline{F}, \bar{F}]$ and $[\underline{V}, \bar{V}]$ respectively are strictly positive.

A *feasible plan* is $(\mathcal{N}^D, \mathcal{N}, X, Y, C)$ Where $\mathcal{N}^D \subset G$ is the choice of research-development projects to undertake, $\mathcal{N} \subset \mathcal{N}^D$ is the choice of goods to actually produce, $Y : \mathcal{I} \rightarrow \Re$ and $X : \mathcal{I} \rightarrow \Re_+^N$ are measurable functions describing the allocation of goods to agents, and $C : \mathcal{I} \times \mathcal{N} \rightarrow \{0, 1\}^N$ is a measurable function that describes whether or not agent i has added good $g \in \mathcal{N}$ to his demand set, such that

$$\int Y(i)di + \sum_{g \in \mathcal{N}} \left(F^g + V^g \int X^g(i)di \right) + t \sum_{g \in \mathcal{N}} \int C(g, i)di + D \times |\mathcal{N}^D| \leq \int \Omega(i)di$$

The material balance condition above says that the amount of the endowment good consumed by agents, plus the amount used to make produced goods, plus the

total transactions costs of adding produced goods to agents' consumption bundles, plus the cost of undertaking $|\mathcal{N}^D|$ development projects is less than or equal to the total social endowment of the non-produced good.

In the following, it will be useful to define the notion of a minimal δ -grid in the innovation space G . Recall that G is bounded. Let \mathcal{N} be a finite subset of G . Then \mathcal{N} is a δ -grid of G if

$$\max_{g \in \mathcal{N}} \min_{\substack{\bar{g} \in \mathcal{N} \\ \bar{g} \neq g}} d(g, \bar{g}) \leq \delta.$$

In other words, a δ -grid is a choice of goods in the innovation space such that each good has at least one neighbor δ or closer under the goods metric d .

Let \mathbf{N}^δ be the set of all δ -grids of G .¹⁴ Then the \mathcal{N} is a *minimal δ -grid of G* if no $\bar{\mathcal{N}} \in \mathbf{N}^\delta$ has a smaller number of elements. Note this is well defined since the magnitude of the set of elements are integers and are bounded below.

We begin by considering how well an individual consumer i could conceivably do in the best of all possible worlds. Thus, we consider the following *idealized consumer problem*: Suppose that an agent could choose any demand set $\mathcal{N}_i \subset G$ and could purchase as much of each good as he wished at price \underline{V} per unit. Formally the idealized consumer consumer problem is:

$$\max_{\mathcal{N}_i \subset G} \max_{\{x_i^g\}_{g \in \mathcal{N}_i}} h_i(\{x_i^g\}_{g \in \mathcal{N}_i}) - tN_i - \underline{V}x^g.$$

¹⁴ If it is not immediately obvious that such a δ -grid exists, consider the following. Note that an open cover of open δ -hyper-balls must exist. To see this, suppose we constructed a sequence of $\frac{1}{2}\delta$ -hyper-balls as follows: Place the first ball anywhere in G . Place the next ball with its center point is at least $\frac{1}{2}\delta$ from every point in the first $\frac{1}{2}\delta$ -hyper-ball. Place the center of each subsequent $\frac{1}{2}\delta$ -hyper-ball such that it is at least $\frac{1}{2}\delta$ from every point in the union of all the previous balls. One of two things is true: Either one can this do forever, or eventually one runs out of space and can not place any more balls. However, if one could do this forever, then one could take the center points of each of these $\frac{1}{2}\delta$ -hyper-balls as a sequence which does not converge under d contradicting the fact that G is compact. If we do eventually run out of space, then every point not contained in the union of $\frac{1}{2}\delta$ -hyper-balls is no more than $\frac{1}{2}\delta$ from some point in one of the existing balls, and thus no more than δ from one of the center points of the existing balls by the triangle inequality. These center points therefore form a finite δ -grid (and of course the implicit set of δ -hyper-balls form a finite open cover of G).

We will call a solution an *idealized consumer demand* and denote it as: \mathcal{N}_i^* and x_i^* respectively. We now show two things about the about this solution.¹⁵

Lemma 1. *There exists an integer \bar{N} such that for all $i \in \mathcal{I}$ and for any solution to the agent's idealized consumer problem, $N_i < \bar{N}$.*

Proof/

By Assumption B, for all $\epsilon_G > 0$ there exists a $\delta_G > 0$ such that for all $i \in \mathcal{I}$ and all integers N for all $\bar{\mathcal{N}}_i = (\bar{g}_1, \dots, \bar{g}_N)$, if for some $\bar{g}_n \in \bar{\mathcal{N}}_i$, $\widehat{\mathcal{N}}_i = \bar{\mathcal{N}}_i \setminus \bar{g}_n$ and for some $\bar{g}_m \in \bar{\mathcal{N}}_i$ $d(\bar{g}_n, \bar{g}_m) < \delta_G$ it is the case that $x_i^{\hat{g}_m} = x_i^{\bar{g}_m} + x_i^{\bar{g}_n}$ while for all $k \neq n, m$, $x_i^{\bar{g}_k} = x_i^{\hat{g}_k}$ then $|U_i(\bar{x}_i, y_i) - U_i(\hat{x}_i, y_i)| < \epsilon_G$. Given this, let $\epsilon_G = t$, that is, the utility cost of adding a good to the demand set, and let δ_G^t be the associated delta. Note that since all goods are equally costly by assumption, cost is not a consideration in deciding of which specific goods to add to the demand set. Then if two goods \bar{g}, \hat{g} are in an agent's demand set but closer than δ_G^t in the characteristic space, then utility is improved by consolidating consumption on one of the goods since the savings from not adding the other goods to demand set exceeds the maximal utility loss from the consolidation. Now consider any minimal δ_G^t -grid. We know that this grid has a finite size which we will denote $\bar{N} - 2$. One of two things is true: First, the size of the demand set chosen in the idealized consumer problem is no more than \bar{N} . In this case, the Lemma is proved. Second, the size of the demand set chosen in the idealized consumer problem is greater than \bar{N} . But note that by construction, there must be an element of the δ_G^t -grid within δ_G^t of any good in the idealized demand set. Then suppose we shifted demand from each good in the idealized demand set to an identical quantity of the closest good in the δ_G^t -grid. Then by Assumption B, utility decreases by at most $\epsilon = t$. However, the agent is consuming at least two fewer goods, so utility must go up by at least $2t$. Since the cost of the two consumption bundles is the same, the agent is strictly better off contradicting that the agent was correctly solving the

¹⁵ The problem is idealized in the sense that the agent can get his most desired set of goods while paying no fixed costs, no development costs and the lowest marginal costs possible. Agents do have to pay the transactions cost of adding goods to their demand set, however.

idealized consumer problem. Thus, the size of any minimal δ_G^t -grid+2 $\leq \bar{N}$ is an upper bound on the number of elements in the consumption set.

■

Lemma 2. *There is an upper bound \bar{U} such that for all $i \in \mathcal{I}$ the utility that any agent can receive at any solution to the agent's idealized consumer problem is less than \bar{U} .*

Proof/

By Assumption C: For all $i \in \mathcal{I}$, h_i is continuous and there exists an upper bound \bar{B} such that for any consumption vector (x_i, y_i) that includes some good g for which $x_i^g > \bar{B}$, it holds that: $\partial h(x_i)/\partial x_i^g < \underline{V}$. Thus, choosing to consume more \bar{B} of any produced good in exchange for \underline{V} per unit of additional consumption would lower an agent's utility. By Lemma 1, there is an upper bound on the number of goods that an agent will have in his demand set. Thus, each agent will consume at most a finite amount of a finite number of goods. As a result, we can restrict the domain of the idealized consumer problems to a compact set bounded above by \bar{B} and bounded below by 0.¹⁶ Since h_i is continuous in x_i , for any choice of \mathcal{N}_i , h_i achieves a finite maximum on this compact set. By Assumption B, U is also continuous in g . Thus, if we fix the number of goods consumed at some N , since G is compact, the joint maximization that chooses exactly N goods and the quantity of each has a solution. By varying N from 0 to \bar{N} , we get a set of solutions. Take the subset of these solutions that give the agent the highest utility to be set of idealized consumer demand. Since by Lemma 1, no solution to the idealized consumer problem in which the consumer is forced to choose more than \bar{N} can yield more utility, we conclude that the solution(s) to the idealized consumer problem must exist, and the Lemma is proved for any single agent. Now consider the set of utilities that agents get at each of these idealized consumer demands. Since the set of all agents is compact, the maximum of these maximum

¹⁶ We assume that agents can't consume negative quantities of the produced good, although they can consume negative quantities of the quasi-linear non-produced good.

utility values is also bounded and so the Lemma is proved for the set of all agents.

■

The two Lemmas above show that there is a theoretical upper bound on the utility that any agent in the economy can obtain and that this involves the agent consuming at most a strictly bounded set of goods. Note, however, that the fact that each consumer chooses only a finite set of goods in no way directly implies that the number of goods demanded by consumers as a whole constitute a finite set. Of course, it may not be feasible to have sufficient product diversity to allow all agents to achieve this maximum, and even if it was, the resulting allocation might or might not be measurable. The next Lemma will show that despite this, there is a plan that allows each agent to get arbitrarily close in expectation to his theoretical maximum utility and which is feasible for a large enough economy.

To see this, we consider how the planner can generate a feasible plan that approximates these idealized utility levels for each agent. Define an (ϵ, R) -*plan* as follows:

1. \mathcal{N}^D : Let $\bar{\mathcal{N}}$ be any minimal δ -grid of G where $\delta = \delta_G$ and δ_G is associated with $\epsilon_G = \epsilon$. Then the set of research and development projects undertaken, \mathcal{N}^D , is defined as developing each project in $\bar{\mathcal{N}}$ a total of R separate times.¹⁷
2. \mathcal{N} : Let $\{g_1, \dots, g_R\}$ be the R separate tries at development in \mathcal{N}^D of any given product in $g \in \bar{\mathcal{N}}$. Then for each $g \in \bar{\mathcal{N}}$ construct \mathcal{N} choosing a $g_i \in \{g_1, \dots, g_R\}$: with the lowest average realization of V^{g_i} and F^{g_i} : $\frac{V^{g_i} + F^{g_i}}{2}$.
3. C: Divide the interval of agents into consecutive coalitions of agents each ϵ units long.¹⁸ For each of these intervals, take the middle agent. Choose any one of the idealized consumer demands for agent i , \mathcal{N}_i^*, x_i^* . For each $g \in \mathcal{N}_i^*$, we will say that good $\bar{g} \in \mathcal{N}$ is in correspondence with good g if it is closest under d :

¹⁷ In other words, take R separate draws from the cost distribution for each good in the δ -grid $\bar{\mathcal{N}}$. If one prefers not to allow reinventing the wheel, one can develop new projects in an arbitrarily small neighborhood of each good in $\bar{\mathcal{N}}$ without affecting the argument in the next Lemma.

¹⁸ To be more precise, the first interval is $[0, \epsilon)$ the second is $[\epsilon, 2\epsilon)$, etc.. If the set of agents divides evenly into ϵ sized intervals, the last interval is closed and includes agent I . Otherwise the last interval is a closed interval of size less than ϵ but that also includes agent I .

$\hat{g} = \operatorname{argmin}_{g \in \mathcal{N}} d(g, \hat{g})$. (In the zero probability event that two goods in \mathcal{N} are the same distance from g choose one of them randomly.) Given this, define \mathcal{N}_i to be the set of goods in \mathcal{N} that are in correspondence with the goods in \mathcal{N}_i^* .¹⁹ Having established the demand set for the middle agent of each interval, we now choose the demand set of every agent j who shares the interval with agent i to be identical: $\mathcal{N}_j = \mathcal{N}_i^*$. Finally, C , the indicator function for which agents are consuming which produced goods, is just constructed to reflect these demand sets over intervals consumers.

4. X: For each agent i in the middle of an interval we set x_i as follows. For every $\hat{g} \in \mathcal{N}$, set $x_i^{\hat{g}}$ to equal to the sum of the consumption levels of each of the goods $g \in \mathcal{N}$ that \hat{g} is in correspondence with. (Again, for a fine grid, the correspondence will be one-to-one). As above, set the consumption levels of each agent j who shares the interval with i to be the same $x_j = x_i$. Finally, X , the produced good consumption mapping, is constructed to reflect these intervals of identical consumption levels.
5. Y: For each agent i in the middle of an interval, set

$$y_i = \omega_i - \frac{RND + \sum_g F^g}{I} - \sum_{g \in \mathcal{N}_i} V^g x_i^g - tN_i.$$

That is, his endowment, minus his average share of the total of the development and fixed production costs in the plan, minus the variable cost of his produced goods consumption, minus the transactions costs of having a consumption set of size N_i . In addition for every agent j who shares the interval with agent i , let

$$y_j = \omega_j - \frac{RND + \sum_g F^g}{I} - \sum_{g \in \mathcal{N}_j} V^g x_j^g - tN_j \equiv \omega_j - \frac{RND + \sum_g F^g}{I} - \sum_{g \in \mathcal{N}_i} V^g x_i^g - tN_i.$$

¹⁹ Intuitively, \mathcal{N}_i chooses a demand set with goods as close as possible to the idealized demand set. If the δ -grid is coarse, then some goods in \mathcal{N}_i chooses a demand set with goods as close as possible to the idealized demand set. If the δ -grid is coarse, then some goods in \mathcal{N}_i^* may not have a close approximation and so they are “doubled-up” with a single good in \mathcal{N}_i . If the grid is fine enough, the correspondence will be one-to-one since we know the goods in the idealized demand set can be too close together by Assumption B.

Thus, the consumption of the non-produced goods for agents in a given interval differs only by the difference in their initial endowment of the good. Finally, Y , the non-produced good consumption mapping, is constructed to reflect these intervals of identical consumption levels.

Lemma 3. *Any (ϵ, R) -plan is feasible.*

Proof/

1. \mathcal{N}^D : As we discussed above, a minimal δ -grid exists as described, and given this, it is clearly feasible to make R development attempts at each good in this grid.
2. \mathcal{N} : Since there are R attempts to develop each product, there are a finite set of pairs of realized (V^{g_i}, F^{g_i}) . Thus, choosing the realization with the lowest average value is well defined (flip a coin in the event of zero probability tie). \mathcal{N} is therefore well defined.
3. C : Dividing the interval of agents into disjoint segments of length ϵ and then choosing the middle agent in each segment is feasible and well-defined. The set of idealized demands is non-empty, and so it is possible to choose one of these and then set each agent in any given interval to have the same demand set as the middle agent. Then $C(i)$ takes values 1 or 0 over the finite set of goods in \mathcal{N} and maintains these same values over the (measurable) ϵ intervals of agents. Thus, $C(i)$ is a simple (step) function and is therefore measurable.
4. X : As above, x_j is constant over each of the ϵ -intervals of agents and thus, it defines a bounded step function which is therefore measurable.
5. Y : For every agent in a given ϵ interval:

$$y_j = \omega_j - \frac{RND + \sum_g F^g}{I} - \sum_{g \in \mathcal{N}_i} V^g x_i^g - tN_i$$

Where i is the middle agent of the interval. By assumption, Ω is a measurable function. Subtracting a constant, $\frac{RND + \sum_g F^g}{I}$ leaves resulting function also measurable. The key point, however, is that $\sum_{g \in \mathcal{N}_i} V^g x_i^g - tN_i$ is also constant over the ϵ interval of agents and changes only across these intervals. Thus, again, it

is a simple (step) function and is therefore measurable. Subtracting a measurable function from a measurable function leaves a measurable function. Thus, the consumption function is measurable.

The last step is to show the plan satisfies material balance. Rearranging this for a generic agent i :

$$y_i + \frac{RND + \sum_g F^g}{I} + \sum_{g \in \mathcal{N}_i} V^g y_i^g + tN_i = \omega_i$$

Integrating over \mathcal{I} we find:

$$\begin{aligned} \int x_i di &= \int X(i) di \\ \int \frac{RND}{I} di &= D \times |\mathcal{N}^D| \\ \int \frac{\sum_g F^g}{I} di &= \sum_{g \in \mathcal{N}} F^g \\ \int \sum_{g \in \mathcal{N}_i} V^g x_i^g di &= \sum_{g \in \mathcal{N}} V^g \int X^g(i) \\ \int tN_i di &= t \sum_{g \in \mathcal{N}} \int C(g, i) di \end{aligned}$$

Thus,

$$\int Y(i) di + \sum_{g \in \mathcal{N}} \left(F^g + V^g \int X^g(i) di \right) + t \sum_{g \in \mathcal{N}} \int C(g, i) di + D \times |\mathcal{N}^D| \leq \int \Omega(i) di$$

■

Our next Lemma show that there is a feasible plan in the class described above that allows each agent to get a utility level that is arbitrarily close to his ideal utility as the economy gets large.

Lemma 4. *For any $\gamma > 0$ there exists a (ϵ, R) -plan such that for a large enough population, each agent in \mathcal{I} receives within γ of the utility he would get at an idealized demand.*

Proof/

Consider an agent i who is in the middle of one of the intervals for some (ϵ, R) -plan. There are three things that might prevent him from getting his idealized utility:

1. It may be that he is forced to consume a set of goods from the δ -grid are slightly less desired than the goods in his idealized demand.
2. It may be that he pays $V^g > \underline{V}$ per unit of some goods he consumes.
3. It will be the case that his share of fixed and development costs is positive and thus he will receive this much less non-produced good, *ceteris paribus*

We consider each in turn. First, if ϵ is small, then the corresponding δ -grid is fine. Moreover, as ϵ heads to zero, goods become available that are arbitrarily close utility substitutes for the goods in his idealized demand set in utility terms by Assumption B. Thus, losses from the first source can be made arbitrarily small by choosing a small enough ϵ . Second, by increasing the number of development efforts for each good in the δ -grid, the variable cost for each good can be driven arbitrarily close to its minimum in expectation since by Assumption D, every outcome has positive probability of occurring. Thus, by increasing R sufficiently, V^g and F^g can be driven as close as one likes to their idealized minimums and thus, per capita losses from this source to below any $\gamma > 0$. Third, having chosen ϵ and R to reduce per capita losses as much as one chooses, fix these and consider an increase in population size I . Since the number of goods is now fixed, the sum of total development and fixed production costs are also fixed. Then, if the population gets large enough, the per capital share can be driven arbitrarily close to zero.

We conclude in the limit all agents will be able to consume an almost ideal set of goods, pay almost the least per unit for the goods he consumes that is physically possible, and also pay almost no per capita fixed and development costs if the economy is large enough.

■

Putting this together, it is almost immediate that the welfare value of product diversity diminish to zero as the number of products increase. In other words, the economy eventually become satiated in product diversity just as agents do in individual products.

Theorem 1. (*Enough is enough*) *The per capita benefits of having greater product diversity diminish to zero and eventually turn negative for any economy.*

Proof/

From Lemma 4, we know that we can get to within any γ of the idealized utility level of every agent for a large enough economy while producing some finite number \bar{N} goods or fewer. Thus, no matter how much the economy continues to grow, at best, increasing the product set adds at most γ per capita utility. On the other hand, for any given economy of fixed size, increasing the number of goods eventually adds more to per capita development and fixed production costs than it does to utility. Thus, increased product diversity will eventually start to lower per capita utility. We conclude that for any economy, enough is enough.

■

We conclude this section by showing the implications of these results for international trade. Consider two parallel identical economies and assume there are iceberg costs of $m \in (0, 1)$ such that only a fraction of m of exported goods arrive at their destination.²⁰

Theorem 2. *It is never per-capita utility improving for a positive fraction of agents to consume any imports as both economies get arbitrarily large.*

²⁰ Given that we have a lower bound the marginal cost of production, this way of specifying iceberg costs implies they are also strictly bounded from zero. Since iceberg costs are supposed to reflect something about the cost of transportation, this seems quite reasonable. Other approaches may allow iceberg costs to approach or even become zero, especially if they are fractional, as they are here, but the marginal cost of production is not bounded from zero. In this case, the Theorem below would not be true, although the per capita benefit of trade would still go to zero. We thank Hideo Konishi for pointing this out.

Proof/

Suppose that a proportion P of the first county consumes a positive amount of an imported good. Then for any $\epsilon > 0$ there would be some good $g \in G$ such that a positive fraction of agents in the first country consumes a positive amount goods that are within an ϵ -ball of g since G is compact. (Of course, it does not have to be the same g 's as the economies grow.) Note that the smallest price that the second country could sell any of these goods in the ϵ -ball of g is $\frac{V}{m}$ since this is the marginal cost when they arrive in the first country. This is strictly larger than \underline{V} . Thus, if enough agents in the first country are consuming goods in an ϵ -ball of g , then it becomes socially optimal to undertake enough development efforts to lower the price below the costs of imports and share the costs of development and the fixed costs of production over the arbitrarily large group that consumes these imports close to g . Since g is a close substitute, is cheaper on the margin, and contributes arbitrarily little to average shares of fixed cost, it would crowd out the imports.

■

4. A Remark on Product Innovation

There are two key modeling assumptions that drive our findings. First, is the inclusion of a tractions cost for adding goods to one's demand set. Economically, this implies that agents will choose at most a strictly bounded set of goods to consume, and certainly will not choose to consume all produced goods if the set of such goods is large. This prevents agents from gaining unbounded utility by the act choosing to consume ever smaller amounts of an ever increasing numbers of goods (taking advantage of the fact that marginal utility of goods goes to infinity as consumption goes to zero as in the Dixit-Stiglitz/Melitz formulations.) This seems both intuitively appealing and clearly borne out by the evidence we see in the real world.

The second assumption is a bit more subtle, but brings up some interesting questions. Note that while we assumed very little structure on innovation space G , we did assume it was compact. This means there is a natural limit on the how well off an agent can be no matter how closely the produced products match his ideal goods (at least under the assumptions that close products are similar and that we have limits of the number of goods we can consume). One can readily see why “enough is enough” in this sort of innovation environment. If the innovation space were unbounded, on the other hand, one could continue invent an infinite number of products that were distant existing ones and so would not have close substitutes. In this case, there could indeed be unbounded per capita benefits from product diversity (although, even in this case, there is no theoretical *requirement* that taste for product diversity are insatiable, not to mention little empirical evidence that this should be the case in the real world). Thus, our assumptions are sufficient to imply product diversity satiation while the assumptions of the DS approach are sufficient to imply non-satiation. In neither case are the assumptions necessary for the conclusions.

Let us therefore consider what taking the bounds off the innovation space might mean. One approach is the product quality ladder type innovation structure with new products potentially being unboundedly good. This has an odd implication, however: in such an environment, agents should get happier and happier as time goes on. We should have been despondent Hittites, miserable subjects of the Roman Empire, unhappy Renaissance men, grumpy Victorians, reasonably content cold warriors and happy Yuppies. Are we measurably happier as a result of the Internet, email, instant messaging and cell phones and Facebook than we were twenty years ago? At best, we did not know what we were missing. At worst, we have to adopt these sometimes intrusive and always-on technologies to keep up with our neighbors. If this seems unlikely going backwards, it seems even more unlikely going forwards. Do we expect that increased product diversity and innovation is really the path to Nirvana?²¹

²¹ We should note that even the quality ladder case, our results would still hold if there was an upper bound on the level of innovation in any given period since this would bound the innovation space each period implying the bounded product diversity was sufficient. See Grossman and Helpman (1991) and references therein.

This leads to a puzzle. If we don't think that we can innovate our way to ultimate bliss how do we explain all the effort that goes into product development and the ready markets that these new products seem to find? What this suggests to us is that the creation of new products may sometimes deprecate the value of existing products. For example, users only realize how lame their current iPhone is when they see the newest model. Up until that point, they imagined that they had the best phone in the world. Now they see it for the sad piece of half-baked technology that it really is. In a formal sense, we prevented this in both of our models by assuming that agents drew utility parameters for new goods once and for all. This discussion, on the other hand, would suggest that some old goods may have their initial draw systematically decreased in the presence of new goods. Of course, in such an environment, firms would profit from developing new goods since optimizing agents would choose to buy these new goods in preference to old ones.

Interestingly, the idea that new goods deprecate old goods does not suggest that increased product innovation and diversity is welfare improving. In fact, agents are no better off when they become disillusioned with their current consumption bundles and substitute into new goods. For example, it may very well be the case that iPhone users would be better off if Apple slowed down the pace new product release. Since consumers are not made better off, innovation in fact may be a waste of resources that uses up a great many talented and creative people who could be otherwise better employed.

Thus, much of technological and product innovation might better be thought of as "fashion". Fashion is an example of innovation for its own sake. We would be perfectly happy to continue wearing this year's fashions and if they were the same next year. However, if a new fashion takes over, we feel silly in our out-of-date clothes and have to spend resources on getting new things. Thus, the change of fashion costs resources both on the development side and on the consumer side as items are replaced before they wear out, but no net improvement in consumer welfare results. On the other hand,

innovations that result in permanent increase in the satisfaction of Maslow’s hierarchy of basic needs²² such as cures for infectious diseases, democratic forms of government, more productive crops, and so on, may not be deprecated by newer innovations.

To summarize: how to model innovation for new products, how agents might evaluate new goods, and how firms should go about estimating demands for goods that don’t currently exist are very difficult questions. Of course, how one deals with these problems from a formal standpoint has large impact on the theoretical outcomes of one’s model. We have proposed a new approach in this paper based on what we called an “innovation space”. While we think this has many advantages, we would certainly not claim that ours is the only or even the best possible approach. We hope, however, that it makes is clear that approach taken in DS based models is not the obvious default. The DS approach to these issues embeds significant assumptions that are non-general, and, we argue, difficult to justify theoretically or empirically. Moreover, the technical assumptions embedded in the DS approach lead directly to conclusions which we think are hard to support.

5. Conclusion

In this paper we have attempted to understand what a Dixit-Stiglitz economy in which the number of products is endogenously determined and potentially unbounded would look like as the population of agents gets large. Our major conclusion is that agents would eventually limit the number of goods they consumed in positive amounts to a strictly bounded subset of all the products offered by firms. This is in contrast to the Dixit-Stiglitz approach which assumes that all agents consume at least some of all goods produced. This in turn forces us to be more explicit about the characteristics of potential goods since agents would have to decide which goods to add to their demand

²² Physiological needs, Safety needs, Love and belonging, Esteem, Self-actualization, Self-transcendence.

sets.

Based on this, we constructed a continuum economy that reflects what we would expect to see in a large finite economy with endogenous product innovation. Our major results are that no matter how large the economy, almost all the benefits of product diversity can be realized by developing a strictly bounded and finite set of goods. In other words, agents become satiated in product diversity. This in turn implies that in the presence of iceberg costs, trade diminishes to zero as economies become large. This is in sharp contrast to the results obtained by Melitz and we have tried to create a convincing case that the standard results must be seen as artifacts of the underlying mathematical structure of the model (in particular, the embedded assumption that all agents consume some of each of the infinite number of goods produced by firms) that are difficult to motivate economically.

At a more technical level, this work suggests that efforts to build tractable algebraic models to understand trade at the cost of making firms into price-takers either by direct assumption or by assuming that there are an infinite number of firms in equilibrium may result in misleading conclusions. In fact, gains from trade (at least when there is no other exogenous heterogeneity over firms) generally come only when there are a small number of products in a sector and so consumers see true benefits from increasing the product mix. This is exactly when firms would not be price-takers.

Our conclusion is therefore that a desire for product diversity is simply not a sufficient foundation upon which one should build a rationalization for trade.

Having said this, trade is clearly important. The model suggests that reasons for trade are actually more classical:

1. Trade should take place when fixed costs are large compared to the size of the economy. Consider aircraft. Fixed costs of development and setting up production lines are so large, that we only have a very few producers of large commercial aircraft. Most countries import such aircraft. We have perhaps a hundred or so car manufactures in the world, but evidently, the five or ten producers of cars in even the largest country are not able to provide sufficient product diversity to

foreclose welfare enhancing trade. Thus, our point is that while a desire from product diversity does in fact motivate trade, this desire is ultimately bounded. Thus, it should only be economically relevant when fixed costs are sufficiently high in comparison to the market demand that only a relatively few firms can enter the product category in equilibrium.

2. Trade should take place when there are differences in abilities. Many French cheese and wine producers are relatively small, and there are many thousands of wine and cheese producers in most large countries. However, we still trade wine and cheese with the French. In the model, this could be explained by the French having a monopoly over part of the innovation space. That is, for technological reasons, there is a region of G that can only be produced in France. The product innovation we model we develop here could easily be generalized to account for this by assuming different, perhaps overlapping innovations spaces for different countries. This means that our model can include Ricardian motivations for trade. In such an environment, if Americans have preferences that make it optimal for some goods g in the region of the innovation space accessible only to the French to be added to their demands sets, these goods should be imported. On the other hand, if American producers could make almost perfect substitutes, it would be irrational to pay the iceberg costs of shipping the French products over.
3. Trade could take place if certain countries are technological leaders. If one takes the view that (a) new innovations deprecate old ones or (b) the boundary of the innovation space G expands each year, then countries who happened to have lower development costs D would have a systematic advantage in bringing out new products. Since less technically apt countries would not innovate as much, there would be a flow new goods from the high-tech countries to the low-tech countries. In contrast, we assumed that all countries were equally able to produce new developments in our model.
4. Trade should take place when there are differences in endowments. China has cheap labor due to its factor endowments. As a result, the assumption that each

country is drawing from the same fixed and variable cost distribution for new innovations is not empirically correct. Even though we have the technology to produce any of the manufactured goods that China sends us, China can produce these at a sufficiently lower price that it can overcome the iceberg costs. Resources may also provide literal monopoly or oligopoly positions in exports. Only a few countries can export diamonds or uranium. Formally, one would model this as certain countries having exclusive abilities to produce goods in certain regions of G . clearly, minor modifications in the our model relating either to access to G or the distribution of costs allow us to include Heckscher-Ohlin motivations for trade as well.

In conclusion we think the predictions of the model we develop here are consistent with real patterns of trade. In particular they suggest that besides standard Ricardian or Heckscher-Ohlin motivations trade, the remaining trade is likely to consist of goods produced by oligopolies. Thus, one must be careful in drawing conclusions from Dixit-Stiglitz based approach which have price-taking behavior and small firms at their foundation.

References

- Balassa, B.** (1967): *Trade liberalization among industrial countries, objectives and alternatives*. Published for the Council of Foreign Relations by McGraw-Hill in New York, London.
- Broda, C., and D. Weinstein** (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economic*, **121** pp. 541-585.
- Dixit, A. K. and Stiglitz, J. E.** (1974): "Monopolistic competition and optimum product diversity," Balliol College, Oxford and Cowles Foundation, Yale,mimeo; see chapter 3 in this volume.
- Dixit A. and J. Stiglitz** (1975): "Monopolistic competition and optimum product diversity," Warwick and Stanford, mimeo; see chapter 4 in this volume.
- Dixit A. and J. Stiglitz** (1977): "Monopolistic Competition and Optimum Product Diversity," *The American Economic Review*, **67** pp. 297-308.

- Grossman, G. and E. Helpman** (1991): "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, **58** pp. 43-61.
- Grubel, H. G. and Lloyd, P. J.** (1975): *Intra-Industry Trade: The Theory and Measurement of International Trade in Differentiated Products*. London, Macmillan.
- Hart, O.** (1985a): "Monopolistic Competition in the Spirit of Chamberlin: A General Model," *Review of Economic Studies*, **52** pp. 529-46.
- Hart, Oliver D** (1985b): "Monopolistic Competition in the Spirit of Chamberlin: Special Results," *Economic Journal*, **95** pp. 889-908.
- Kaneko, M. and Wooders, M.H.** (1986): "The core of a game with a continuum of consumers and finite coalitions: The model and some results," *Mathematical Social Sciences*, **12** pp. 105-137.
- Krugman, P. R.** (1979): "Increasing returns, monopolistic competition and international trade," *Journal of International Economics*, **9** pp. 469-479.
- Krugman, P.** (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *The American Economic Review*, **70** pp. 950-959.
- Krugman, P. R. and M Obstfeld** (2009): *International economics: theory and policy*. Pearson Addison-Wesley.
- Lancaster, K. J.** (1980): "Intraindustry trade under perfect monopolistic competition," *Journal of International Economics*, **10** pp. 151-175.
- Mas-Colell, A.** (1980): "Efficiency and Decentralization in the Pure Theory of Public Goods," *Quarterly Journal of Economics*, **47** pp. 239-56.
- Melitz, M.** (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, **71** pp. 1695-1725.
- Melitz, Marc and Ottaviano, G.** (2008): "Market Size, Trade, and Productivity," *The Review of Economic Studies*, **75** pp. 295-316.
- Mrazova, M. and P. Neary** (2011): "Selection effects with heterogeneous firms," Economics Series Working Papers 588, University of Oxford, Department of Economics.
- Neary, J. Peter** (2004): "Monopolistic competition in international trade theory," in *Monopolistic Competition Revolution in Retrospect*, by Brakman, Steven, and Ben J. Heijdra, editors. Cambridge University Press, eBook Collection (EBSCOhost): p. 159-84/
- Perloff, J. and S. Salop** (1985): "Equilibrium with Product Differentiation," *Review of Economic Studies*, **52** pp. 107-20.