Business Networks, Production Chains, and Productivity*

Ezra Oberfield

University of Chicago and Federal Reserve Bank of Chicago

ezraoberfield@gmail.com

February 9, 2011

Abstract

This paper studies an analytically tractable model of the formation and evolution of chains of production. Over time, entrepreneurs accumulate techniques to produce a good using labor and some other good as an intermediate input. The value of a technique depends on both the productivity embodied in the technique and the cost of the particular inputs. When producing, each entrepreneur selects the technique that delivers the most cost effective combination. The collection of known production techniques form a web of interweaving and overlapping potential chains of production: the input-output architecture of the economy. The model delivers a rich microstructure of firm level shocks and interactions as firms discover new techniques and switch suppliers in response to changes in input prices. Changes in a firm's marginal cost depend on shocks all along its many potential supply chains, and the size of a firm's customer base expands as other entrepreneurs find ways to use the firm's output as an intermediate input. Dispersion across firms in the rate which links are formed can have enormous consequences for aggregate output (either positive or negative), as this determines how productivity gains from newly discovered techniques diffuse through the network.

*Preliminary and Incomplete, please do not quote without permission of the author. I appreciate the comments of Amanda Agan, Fernando Alvarez, Gadi Barlevy, Marco Bassetto, Jeff Campbell, Thomas Chaney, Aspen Gorry, Sam Kortum, Alejandro Justiniano, Robert Lucas, Devesh Raval, Rob Shimer, Nancy Stokey, Nico Trachter, and Andy Zuppann, as well as seminar participants at the University of Chicago and the Federal Reserve Bank of Chicago. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Chicago or the Federal Reserve System. All mistakes are my own.

An entrepreneur is constantly searching for new ways to produce. When he meets another entrepreneur, he may develop a new technique for making his good using the other's as an input. Over time he accumulates different ways of producing his good. If one of his potential suppliers innovates and is able to lower her price, that technique becomes more attractive. One can visualize the set of all techniques known to entrepreneurs as a network, where each node is an entrepreneur and techniques are links between them. When any of the entrepreneurs discovers a new, more efficient technique, the productivity gains diffuse through the network. How large those effects are and how quickly these more efficient techniques are developed hinge on both the structure and the density of the network.

When many entrepreneurs are in close proximity and are exposed to each others' products, techniques are discovered frequently, giving each entrepreneur many options to choose from. At a point in time, each entrepreneur uses only her most cost effective technique, but as an economy evolves substitution across techniques can be very important. For example, a technique that uses oil to produce electricity might be useful when oil is cheap; when the price of oil rises, the electricity producer may substitute toward a technique that uses coal.

Over time, the collection of known production techniques form a dynamic web of interweaving and overlapping potential chains of production. One can think of the information stored in this network as the set of supply chains (chains of techniques) available to make each final product. Changes in a entrepreneur's marginal cost can come from changes anywhere along her many supply chains, and these are passed on to the entrepreneur's customers. A technique that is not currently cost effective enough may later become so if there are enough cost reductions upstream.

Aggregate productivity is literally embedded in the network of techniques, and both the density and the structure of the network matter. In a more dense network, each entrepreneur has access to more techniques which, on average, allow them to find more efficient production methods. This allows them to lower their costs, passing those gains on to their customers, further increasing aggregate productivity.

The structure of the network determines how broadly and quickly productivity gains diffuse to other entrepreneurs through lower prices. Consider an entrepreneur who discovers an extremely cost effective technique. If the firm has many potential customers (i.e., many other firms have access to techniques that use the good), those lower prices will ripple down many supply chains. If, however, the firm has few potential customers, the lower prices may not have much of an immediate impact. Eventually, if other firms are able to discover techniques that use the good, the innovation will become socially beneficial.

An advantage of the setup is its tractability. For the baseline case, there is a closed form expression for aggregate output, and more generally a simple formula relating the density of the network to aggregate productivity. A simple extension that allows for alternative configurations of the network, with the finding that dispersion in the rate at which firms are involved with newly formed techniques can have enormous consequences for aggregate productivity, both positive and negative.

The notion that business networks play a large role in economic activity has a long history. For an entrepreneur to fully take advantage of her talents, it is useful to have many contacts: to know which other firms would make the best business partners, to get access to better fitting and cheaper inputs, and to find other entrepreneurs who can make use of the good she produces. While such interactions are no doubt an important part of economic activity, it can be challenging to write down a model that is tractable enough to get at the aggregate implications but also flexible enough to begin to quantify the importance of these kinds of interactions. The goal of this paper is to take a step in this direction. The main contribution of the paper is finding a probabilistic structure that allows the incorporation of such a network into a macro model in a tractable way.

Understanding how business networks and chains of productions impact productivity can be useful in several areas. For example, in a dense city entrepreneurs come across each other frequently, potentially given them ideas for new techniques. Some of the gains from agglomeration within a city may come from putting entrepreneurs in close proximity. This gives rise to a more densely populated network of techniques in which firms have more options to choose from, leading to more efficient production and higher aggregate output.[1]

Alternatively, consider a small village in which the state of technology is poor. Transferring several techniques may not be helpful as a technique has no value in isolation; without the required inputs the technique cannot be used. Further, if the productivity embodied in the technique is relationship specific (e.g., two entrepreneurs work particularly well together), then transferring the techniques may not even be feasible.

This paper is related to several disparate literatures. While there are other models that use networks, the structure of the model is most closely connected to the work of Kortum (1997), Eaton and Kortum (2002), Alvarez et al. (2008), and Lucas (2009) who study flows of ideas.

The idea that network structure determines how shocks propagate through an economy has arisen in the real business cycle literature, beginning with Long and Plosser (1983). Recently, the discussion has centered on whether shocks to particularly well connected sectors can prove important enough to account for aggregate fluctuations.[2] These models typically assume that each

[1] This complements but is a distinct channel from the learning in Lucas (2009) in which proximity makes it easier for good ideas to be copied.

[2] See Horvath (1998), Dupor (1999), Carvalho (2007), Acemoglu et al. (2010), and Foerster et al. (2008). This literature has focused on the sectoral level both because the most fine input-output data is at that level and because solving these models involves inverting matrices, which becomes computationally intensive as the number of nodes in the network grows large.

sector has a representative firm that produces using a Cobb-Douglas production function, using inputs from each other sector.

While there are several differences between those models and the one presented here, let me point to an interesting connection. The model presented here provides a microfoundation for the Long and Plosser model in the following sense: If firms are divided into sectors, and all techniques in a sector are augmented by a sector specific neutral technology shock, then a log-linear approximation around the steady state is the Long and Plosser model. In this approximation, there is directed mapping between the share parameters of the Cobb-Douglas sectoral production function in the Long and Plosser model and the quantity of techniques that firms in one sector have formed that use inputs from firms in a different sector in this model. While the model responds similarly to sectoral productivity shocks when aggregated to the sectoral levels, this masks a lot of activity under the hood. In this model, in response to a negative productivity shock to all firms in a sector, some of those firms' customers will choose alternative supply chains in response to the higher input prices, while other customers will simply swallow the higher prices and pass those on to *their* customers.

Models in this literature have taken the sectoral input-output structure as a primitive, with the explicit assumption that the representative firms at the sectoral level have fixed Cobb-Douglas production functions. This is difficult to reconcile with the large long-run changes in input-output shares over time. The particular microfoundation provided by this model is particularly attractive in that it allows for the Long and Plosser model as a local approximation but also gives a clean interpretation of changes in the sectoral input-output matrix over time.

Jones (2008) uses a similar model to argue the input-output structure can help us understand cross country income differences. In that setup misallocation in one sector raised input prices in other sectors, and the magnitude of the overall effect depends on the input-output structure.

This paper proceeds as follows: Section 1 describes the basic technology, setting up and solving a social planner's problem. In this section there are simple formulas relating the density of the network to aggregate output. In addition, I discuss market structures that could be overlaid on the economic environment. Section 2 describes the size distribution. Section 3 generalizes the model presented in the first section to multiple types, to allow for more interesting network configurations. Section 4 describes how alternative configurations relate both aggregate output and the way productivity gains from new techniques diffuse through the network.

1 The Baseline Model

1.1 Economic Environment

There is a mass of infinitely-lived firms, and each firm is associated with a particular good. Each good can be used both as an intermediate input and for final consumption. A representative consumer has Dixit-Stiglitz preferences over the goods and supplies labor inelastically (both of these can easily be relaxed).

At a given point in time, firm j has access to several techniques to produce its good. Each technique $\phi = \{j, i, z\}$ consists of three components: (i) the good that is produced, j; (ii) the good used as an input, i; and (iii) a production technology associated with using that input, indexed by the productivity parameter z:

$$y = \frac{1}{\alpha^{\alpha}(1-\alpha)^{1-\alpha}} z x^{\alpha} L^{1-\alpha}$$

where y is the quantity of output of good j produced, x is the quantity of good i used as an input, L is labor.

Given a menu of techniques and input prices, each firm will maximize profits. In principle a firm can produce using more than one technology, but since the production functions are constant returns to scale, generically the firm will choose the best price/productivity combination available. At a given time, the state of the economy can be summarized by the set of available techniques, $\Phi(t)$.

Firms discover and lose access to techniques (draws) over time according to a random process. Time is continuous and each firm receives draws at the arrival rate $\tilde{\Lambda}(t)$. A technique that exists becomes infeasible at rate δ.[3]

Let $\omega(n,t)$ be the fraction of firms with access to n techniques at time t. It will be useful to characterize the evolution of $\omega(n,t)$ over time. This satisfies:

$$\dot{\omega}(n,t) = \tilde{\Lambda}(t)\omega(n-1,t) + (n+1)\delta\omega(n+1,t) - \tilde{\Lambda}(t)\omega(n,t) - n\delta\omega(n,t) \tag{1}$$

$\omega(n,t)$ increases when a firm with $n-1$ techniques discovers a new one and when a firm with $n+1$ techniques loses one of their $n+1$ techniques. Similarly $\omega(n,t)$ decrease when a firm with n techniques either gains a new technique or loses one of its n techniques. If at some t_0 the distribution of $\omega(n,t_0)$ is given by a poisson distribution with mean $\tilde{\lambda}(t_0)$, then a solution to equation (1) is such that $\omega(n,t)$ is also given by a Poisson distribution[4] with mean $\tilde{\lambda}(t)$, where $\tilde{\lambda}(t)$ satisfies the differential equation

$$\dot{\tilde{\lambda}}(t) = \tilde{\Lambda}(t) - \delta\tilde{\lambda}(t) \tag{2}$$

To interpret this, it helps to take the limit as $t_0 \to -\infty$ (and imposing that $\tilde{\lambda}(t_0)$ doesn't blow up

[3] δ plays a small role in the analysis and setting $\delta = 0$ would change little. It is included (i) for generality and (ii) so that when $\tilde{\Lambda}(t)$ is constant there is a well defined steady state.

[4] In fact, for an arbitrary initial distribution, the distribution of links will converge asymptotically to Poisson distribution.

in the process) giving

$$\tilde{\lambda}(t) = \int_{-\infty}^{t} e^{-(t-\tau)\delta} \tilde{\Lambda}(\tau) d\tau$$

This is closely related to the fact that the sum of independent Poisson random variables is also a Poisson random variable. The takeaway from this is that regardless of history of $\left\{\tilde{\Lambda}(\tau)\right\}_{\tau \leq t}$, the distribution of $\omega(n,t)$ at a given point in time can be summarized by a single number, $\tilde{\lambda}(t)$.

1.2 Planner's Problem

For the most part, I will focus on a planner's problem in order to build intuition about the economic environment and to describe solution techniques without getting bogged down with the details of a particular market structure. There are many market structures that could be layered on top of the technological environment. In Section 1.8 I will discuss two such market structures, both of which decentralize the planner's solution.

Consider the problem of a planner that takes the network of techniques as given but seeks to make production decisions and allocate labor to maximize the utility of the representative agent. Let y_j^0 be production of good j for final consumption, and let Φ_j be the set of techniques available to produce good j. For a technique $\phi = \{j, i, z\}$, define the following quantities:

- $y_j(\phi)$ is the quantity of good j produced using technique ϕ
- $x_i(\phi)$ is the quantity of good i used as an input in the production of j using technique ϕ
- $L(\phi)$ is the quantity of labor used in production of j using technique ϕ
- $z_{ji}(\phi)$ is the productivity parameter associated with technique ϕ

Formally, the planner chooses an allocation $\left\{y_j^0, \{y_j(\phi), x_i(\phi), L(\phi)\}_{\phi \in \Phi_j}\right\}_j$ to maximize final consumption:

$$\max \left(\int_{j \in J} (y_j^0)^{\frac{\varepsilon}{\varepsilon-1}} \right)^{\frac{\varepsilon}{\varepsilon-1}}$$

subject to: (i) technological constraints

$$y_j(\phi) \leq \frac{1}{\alpha^\alpha (1-\alpha)^{1-\alpha}} z_{ji}(\phi) x_i(\phi)^\alpha L(\phi)^{1-\alpha}, \qquad \forall \phi \in \Phi_j, j \in J$$

(ii) goods feasibility constraints

$$y_j^0 + \sum_{\phi \in \Phi_k, k \in J} x_j(\phi) \leq \sum_{\phi \in \Phi_j} y_j(\phi), \qquad \forall j \in J$$

and (iii) a labor resource constraint

$$\int_{j \in J} \sum_{\phi \in \Phi_j} L(\phi) \leq L$$

On the left hand side of the second constraint consists is the sum of final output of good j and output used as an intermediate among all of the techniques that use good j. On the right hand side is the quantity of good j produced using each of the techniques available. In the third constraint, the left hand side is the total labor used with each technique across all firms.

Let MC_j be the marginal social cost of producing good j (the multiplier on the goods feasibility constraint for j), and let w be the marginal social cost of labor (the multiplier on the labor resource constraint). The first order necessary conditions from this problem imply that for each $\phi \in \Phi_j$,

$$\frac{MC_j}{w} \leq \frac{1}{z_{ji}(\phi)} \left(\frac{MC_i}{w} \right)^\alpha \qquad \text{(with equality if } y_j(\phi) > 0) \tag{3}$$

For each technique, the right hand side of equation (3) gives the marginal social cost of producing good j using that technique. The left hand side is the actual marginal social cost of producing good j, i.e., the marginal social cost associated with the technique that is actually used.

It will be convenient to define $q_j \equiv \frac{1}{MC_j}$ as a measure of the efficiency of producing good j. If we choose units of utility so that $w = 1$, we can rewrite equation (3) as

$$q_j \geq z_{ji}(\phi) q_i^{\alpha} \qquad \text{(with equality if } y_j(\phi) > 0) \tag{4}$$

1.3 The Cross Sectional Distribution

Let $F(q)$ be the cross section distribution of efficiency given the decision of the planner. This is an object that will need to be solved for.

If firm j gets a single draw of a potential supplier, there are two parts that determine how useful it is: a productivity parameter, z, drawn from an exogenous distribution $H(z)$, and the efficiency of the supplier, q_i. Recall from equation (4) that if firm j use a positive quantity of input i, firm j will produce at efficiency $q_j = z_{ji}(\phi) q_i^{\alpha}$. Let $G(q)$ be the cumulative distribution of the efficiency associated with a single random draw. Given equation (4), we can write $G(q)$ as

$$G(q) = \int_0^{\infty} F\left(\left(\frac{q}{z}\right)^{\frac{1}{\alpha}}\right) dH(z) \tag{5}$$

To interpret this, note that for each z, $F\left(\left(\frac{q}{z}\right)^{\frac{1}{\alpha}}\right)$ is the portion of potential suppliers that, in combination with that z, would leave the firm with efficiency less than or equal to q.

We now ask, what is the probability that, given all of its draws, a firm has efficiency less than

q? We can write this as

$$\begin{aligned}
\Pr(Q \leq q) &= \sum_{n=0}^{\infty} \Pr(\text{All } n \text{ draws are} \leq q)\,\omega(n) \\
&= \sum_{n=0}^{\infty} G(q)^n \frac{\tilde{\lambda}^n e^{-\tilde{\lambda}}}{n!} \\
&= e^{-\tilde{\lambda}[1-G(q)]}
\end{aligned}$$

To interpret this last expression, if $\tilde{\lambda}[1-G(q)]$ is a parameter of a Poisson distribution (the arrival rate of techniques that would provide efficiency better than q), then $e^{-\tilde{\lambda}[1-G(q)]}$ is the probability that no such techniques arrived.

When the number of firms is large, a standard abuse of the law of large numbers gives $\Pr(Q \leq q) = F(q)$. We can substitute the expression for $G(q)$ from equation (5) to get a fixed point problem for the distribution of efficiency $F(q)$:[5]

$$F(q) = e^{-\tilde{\lambda}\left[1-\int_0^\infty F\left(\left(\frac{q}{z}\right)^{\frac{1}{\alpha}}\right)dH(z)\right]} \tag{6}$$

This is a key equation, highlighting the recursive nature of the network. The behavior of network depends on whether $\tilde{\lambda} \gtrless 1$. The more interesting case in which $\tilde{\lambda} \gg 1$ will be the focus of the remainder of this paper. For completeness I will discuss both cases here, but one could skip to the next subsection.

Few Techniques: $\tilde{\lambda} \leq 1$

If firm j does not have access to any techniques, it cannot produce. Similarly, if firm j has

[5] We could similarly write this as an equivalent fixed point problem in for $G(q)$, the distribution of efficiency provided by a single technique:

$$G(q) = \int_0^\infty e^{-\tilde{\lambda}\left[1-G\left(\left(\frac{q}{z}\right)^{\frac{1}{\alpha}}\right)\right]}dH(z)$$

techniques but its suppliers do not, then those suppliers will not be able to produce and consequently neither will firm j. Continuing with this logic, if a supply chain is finite, it is not viable.

When $\tilde{\lambda} \leq 1$, there are so few techniques available almost all firms have no access to viable supply chains. For almost all of the goods the marginal social cost of producing is infinite. Consequently, $F(q) = 1$ for all q. When $\tilde{\lambda} \leq 1$, one can show that the mapping $f \mapsto e^{-\tilde{\lambda}\left[1-\int_0^\infty f\left(\left(\frac{q}{z}\right)^{\frac{1}{\alpha}}\right) dH(z)\right]}$ is a contraction with the unique solution $f = 1$.

Many Techniques: $\tilde{\lambda} > 1$

As $\tilde{\lambda}$ crosses the critical value of 1, enough techniques are added that a positive fraction of firms will be able to produce. One can show that given $\tilde{\lambda}$, the fraction $\underline{F}$ will have no access to viable supply chains and will be unable to produce. $\underline{F}$ is the minimum solution to the equation $F = e^{-\tilde{\lambda}(1-F)}$, and specifically if $\tilde{\lambda} > 1$ then $\underline{F} \in (0, 1)$.

When $\tilde{\lambda} > 1$, there are multiple solutions to equation (6). Note that since equation (6) was derived from necessary conditions, not sufficient conditions, so one must check which of these solutions to equation (6) actually solves the planner's problem.[6]

First note that there are two solutions in which $F(q)$ is constant for all q. The first is $F(q) = 1$ for all q, which again corresponds to zero efficiency (infinite marginal social cost) for all goods. The rationale is different than in the $\tilde{\lambda} \leq 1$ case; here, if the marginal social cost of every input is infinite, then the marginal social cost of each output must be infinite as well. The allocation that arises from this solution is feasible, but it is dominated by another feasible allocation and is therefore not the solution to the planner's problem.

There is a second constant solution, $\underline{F} \in (0, 1)$, which follows a similar circular logic. The

[6] As will be clear from the logic, multiple solutions to necessary conditions are actually a feature of most models in which a portion of output is used simultaneously as input, such as a standard growth model with roundabout production.

constant solution $F(q) = \underline{F}$ corresponds to infinite marginal social cost for those firms that cannot produce, and zero marginal social cost for all other firms. The rationale is similar: if inputs have zero marginal social cost, output has zero social cost. Unfortunately, this leads to an infeasible allocation, and is therefore not a solution to the planners problem either.[7]

There exists a third solution as well, which is feasible and solves the social planners problem. This will be discussed further below.

The starkly different behavior of the network when $\tilde{\lambda}$ crosses 1 is called a phase transition. This phase transition is a typical feature of random networks, a result known as the Erdos-Renyi Theorem.[8]

1.4 The Allocation of Labor and Welfare

Once we have solved for the cross sectional distribution of efficiency given the network of techniques, we next need to assign the correct quantity of labor to each firm to ensure sufficient quantities of intermediate goods are produced. We take advantage of a convenient property of the model. Let L_j^k be the quantity of labor used in the kth-to-last stage of production of good j (so the L_j^0 is the labor used in production of final good j, L_j^1 is the labor used to make the inputs for that, etc.). Further, define $L^k \equiv \int_J L_j^k$ to be the total labor used for the kth-to-last stage of production across

[7] These two constant solutions have further economic meaning. Given the distribution of productivity draws $H(\cdot)$ with support $[\underline{z}, \bar{z}]$ with $0 \leq \underline{z} \leq \bar{z} \leq \infty$, let $\underline{q}$ and $\bar{q}$ be the lowest and highest possible efficiencies among firms that are able to produce. $\underline{q}$ ($\bar{q}$) is derived from the supply chain in which every technique has the worst (best) possible productivity draw, so that $\underline{q} = \underline{z}^{\frac{1}{1-\alpha}}$ ($\bar{q} = \bar{z}^{\frac{1}{1-\alpha}}$). If $\tilde{\lambda} > 1$ the solution to the planners problem must have $F(\bar{q}) = 1$ and $F(\underline{q}) = \underline{F}$, the two constant solutions to equation (6). One can see that it must be the case that $F(\bar{q})$ and $F(\underline{q})$ are constant solutions to equation (6) by thinking through the derivation of that equation.

[8] See Kelly (1997) and Kelly (2005) for examples in which this kind of phase transition is given an economic interpretation.

all goods. It is straightforward to show that

$$L_j^{k+1} = \alpha L_j^k \tag{7}$$

and by aggregating across firms it easily follows that $L^{k+1} = \alpha L^k$. So, total labor used in the economy across all stages of production, is

$$L = \sum_{k=0}^{\infty} L^k = \sum_{k=0}^{\infty} \alpha^k L^0 = \frac{1}{1-\alpha} L^0 \tag{8}$$

If L units of labor are supplied, then $L^0 = (1-\alpha)L$ units are used in the final stage of production.

We now use several more first order conditions from the planners problem to arrive at an expression for total final consumption. Define $Q \equiv \left(\int_J q_j^{\varepsilon-1}\right)^{\frac{1}{\varepsilon-1}}$, a standard productivity aggregator for economies with Dixit-Stiglitz preferences. The first order conditions with respect to y^0 and $L_(\phi)$ imply $\frac{y_j^0}{Y^0} = \left(\frac{q_j}{Q}\right)^{\varepsilon}$ and $(1-\alpha)\frac{y_j(\phi)}{L(\phi)} = wq_j$ respectively for each j. Given the constant returns to scale in production, the latter equation implies $y_j^0 = \frac{1}{1-\alpha} wq_j L_j^0$. Combining these gives (recalling that we chose units of utility so that $w = 1$)

$$\frac{L_j^0}{1-\alpha} = Y^0 Q^{-\varepsilon} q_j^{\varepsilon-1}$$

We can now use the labor resource constraint to write

$$L = \frac{L^0}{1-\alpha} = \int_J \frac{L_j^0}{1-\alpha} = \int_J Y^0 Q^{-\varepsilon} q_j^{\varepsilon-1} = Y^0/Q \tag{9}$$

or more conveniently

$$Y^0 = QL \tag{10}$$

1.5 The Supply Chain Interpretation

Given the structure of the network, we can back out the production technology used by the social planner to produce each good. With notation analogous to Section 1.4, we can define q_j^k to be the efficiency of the kth firm in the chain of production for good j, and z_j^k to be the productivity parameter in production of the kth step. This means that along the supply chain for good j, $q_j^k = z_j^k \left(q_j^{k+1}\right)^\alpha$. By definition, $q_j^0 = q_j$, so we make repeated substitutions to get

$$\begin{aligned} q_j &= q_j^0 = z_j^0 \left(q_j^1\right)^\alpha = z_j^0 \left[z_j^1 \left(q_j^2\right)^\alpha\right]^\alpha = ... \\ &= \prod_{k=0}^{\infty} \left(z_j^k\right)^{\alpha^k} \end{aligned}$$

From above, we know that the total quantity of labor used to make j across all stages of production is $\bar{L}_j \equiv \frac{1}{1-\alpha} L_j^0$. For the final stage of production, we know that $y_j^0 = \frac{1}{1-\alpha} w q_j L_j^0$, so substituting in, we have an expression for a production function describing the total social cost of producing good j:

$$y_j^0 = \left[\prod_{k=0}^{\infty} \left(z_j^k\right)^{\alpha^k}\right] \bar{L}_j \tag{11}$$

1.6 A Parametric Assumption

I now describe a special case that proves to be analytically tractable. Assume that the productivity parameter embodied in a technique is drawn from a Pareto distribution, $H(z) = 1 - \left(\frac{z}{z_0}\right)^{-\zeta}$. We make the restriction that $\zeta > \varepsilon - 1$ so that utility will be bounded. In addition, parameterize the arrival rate (and initial condition) of draws so that $\tilde{\Lambda}(t) = \Lambda(t) z_0^{-\zeta}$ which implies that $\tilde{\lambda}(t) = \lambda(t) z_0^{-\zeta}$.

With these assumptions, lowering z_0 has two effects: (i) each firm discovers new techniques more frequently and (ii) a larger fraction of the techniques have low productivity. In fact, the parameterization is such that varying z_0 has no impact on the number of draws above any threshold $\hat{z}$ (above z_0), the two effects cancel exactly. The only change that occurs with a lower z_0 is that there are now more (extra) relatively bad ideas.

We then look at the limit of a sequence of economies as $z_0 \to 0$. This adds many relatively unproductive techniques (low z) to the economy without changing the number of productive techniques (high z). The assumption ensures that the measure of firms without access to any techniques goes to zero, filling out more of the network.

In this special case, we can show that every solution $F(\cdot)$ to equation (6) follows a Frechet distribution. To see this, note that we can use the change of variables $x = (q/z)^{1/\alpha}$ to write

$$\begin{aligned} 1 - G(q) &= \int_{z_0}^{\infty} H'(z) \left(1 - F\left(\left(\frac{q}{z}\right)^{\frac{1}{\alpha}}\right)\right) dz \\ &= \int_{0}^{\left(\frac{q}{z_0}\right)^{\frac{1}{\alpha}}} H'\left(\frac{q}{x^{\alpha}}\right) (1 - F(x))\, q\alpha x^{-\alpha-1} dx \end{aligned}$$

Using the functional form $H'(z) = \zeta z_0^{\zeta} z^{-\zeta-1}$, we can then write

$$\begin{aligned}
\tilde{\lambda}\left[1 - G(q)\right] &= \lambda z_0^{-\zeta} \int_0^{\left(\frac{q}{z_0}\right)^{\frac{1}{\alpha}}} \zeta z_0^{\zeta} \left(\frac{q}{x^{\alpha}}\right)^{-\zeta-1} \left(1 - F(x)\right) q \alpha x^{-\alpha-1} dx \\
&= q^{-\zeta} \lambda \int_0^{\left(\frac{q}{z_0}\right)^{\frac{1}{\alpha}}} \alpha \zeta x^{\alpha\zeta-1} \left(1 - F(x)\right) dx
\end{aligned}$$

For any $F(\cdot)$, as $z_0 \to 0$, this expression will clearly go to $q^{-\zeta}$ multiplied by a constant. Label this constant θ, so that equation (6) can be written as $F(q) = e^{-\theta q^{-\zeta}}$, the cumulative distribution of a Frechet random variable. Note that the exponent ζ is the same as that of the Pareto distribution H.

We next solve for θ, which was defined to satisfy

$$\theta = \lambda \int_0^{\infty} \alpha \zeta x^{\alpha\zeta-1} \left(1 - F(x)\right) dx$$

Integrating by parts gives

$$\theta = \lambda \int_0^{\infty} x^{\alpha\zeta} F'(x) dx \tag{12}$$

Plugging in the functional form $F(q) = e^{-\theta q^{-\zeta}}$ and making the substitution $s = \theta x^{-\zeta}$ gives

$$\theta = \lambda \int_0^{\infty} \theta^{\alpha} s^{-\alpha} e^{-s} ds$$

so that θ satisfies

$$\theta = \Gamma(1-\alpha) \lambda \theta^{\alpha}$$

where $\Gamma(\cdot)$ is the gamma function.[9]

With this, we can compute Q, the relevant measure of welfare:

$$Q = \left(\int_0^\infty q^{\varepsilon-1} dF(q)\right)^{\frac{1}{\varepsilon-1}} = \theta^{1/\zeta}\Gamma\left(1 - \frac{\varepsilon-1}{\zeta}\right)^{\frac{1}{\varepsilon-1}}$$

Putting these together, we get an expression for final consumption:

$$Y = \kappa\lambda^{\frac{1}{1-\alpha}\frac{1}{\zeta}} L \tag{13}$$

where the constant of proportionality is $\kappa = \Gamma(1-\alpha)^{\frac{1}{1-\alpha}\frac{1}{\zeta}}\Gamma\left(1 - \frac{\varepsilon-1}{\zeta}\right)^{\frac{1}{\varepsilon-1}}$.

1.7 Dynamics

All relevant aggregate dynamics can be summarized by two equations

$$Y = \kappa\lambda^{\frac{1}{1-\alpha}\frac{1}{\zeta}} L$$

and

$$\dot{\lambda}(t) = \Lambda(t) - \delta\lambda(t)$$

If the (normalized) arrival rate of new techniques $\Lambda(t)$ is constant over time (and if $\delta > 0$), then there is a steady state with

$$\lambda_{ss} = \frac{\Lambda}{\delta}$$

Alternatively if the arrival rate of new techniques is growing over time, say $\Lambda(t) = \Lambda e^{\gamma t}$, then

[9] I write the equation this way rather than solving for θ directly in order to emphasize the fact that the equation has three non-negative roots, two of which are zero and infinity. In addition, in later sections it will be easier to see parallels with the analogous expressions when this equation when written in this form.

there is a balanced growth path (for any $\delta \geq 0$) with

$$\frac{\dot{\lambda}}{\lambda} = \gamma \qquad \text{and} \qquad \frac{\dot{Y}}{Y} = \frac{1}{1-\alpha}\frac{1}{\zeta}\gamma$$

1.8 Market Structure and Decentralized Equilibria

TO BE COMPLETED

There are at least two valid interpretations of the technology, and for each interpretation there is a natural market structure. The first interpretation is basically the one described in the introduction. Individual entrepreneurs discover techniques, and the technology embodied in each technique is relationship specific. When relationship specific productivity is important, there is an element of bilateral market power. In this kind of environment, a market structure that involves might be a natural benchmark.

An alternative interpretation is that the technology embodied in a technique is non-rival and freely available for others to replicate. In this interpretation, each good is produced by an island of identical firms, and labor is perfectly mobile across islands. The state of technology can be still be represented as a network, but each node is an island of firms producing a single good rather than an individual entrepreneur. Again, the network represents the input-output architecture of the economy, but among islands of firms. In this interpretation perfect competition might be a more natural benchmark with all prices are set at marginal cost. It is straightforward to show that the allocation will decentralize the planner's solution.

In the first interpretation with relationship specific technology, the problem is trickier. I will argue that the most natural interpretation is to allow for two part tariffs, and that this will also decentralize the planners solution.

Monopolistic competition across final goods leads to a uniform markup of $\frac{\varepsilon}{\varepsilon-1}$. This gives surplus for each final good that is divided across firms in that supply chain. If firms are allowed to bargain bilaterally and write complete contracts, then I conjecture that the resulting equilibrium will decentralize the efficient allocation. Contracts will resemble two part tariffs, with a fee and a per unit price. The per unit price will equal marginal cost, but firms will bargain over the fee to split the surplus from the relationship.

In this decentralization, the inputs, outputs, and labor used by each firm are the same as the the efficient allocation. Each firm will choose the supplier that gives the best combination of cost/productivity. While the division of surplus can be difficult to characterize, it has no effect on the allocation of goods or labor. Since the marginal input price equals marginal cost, the quantity supplied will be efficient, and since labor is supplied inelastically the monopoly markup on final goods is not distortionary[10].

This decentralization seems reasonable because input output relationships are generally long term, so it would be surprising if the contracting terms remained inefficient. Anything other than this kind of two part tariff would lead to double marginalization and leave surplus on the table. In addition, the informational demands of these complete contract (in equilibrium) are not large. Firms do not need to know that much about the environment to get the terms of the contract right.

I also believe this is the only tractable way to think about a decentralized equilibrium. Demand curves facing firms are not continuous let alone differentiable, as lowering a price a little may allow a supplier to beat out a competitor and give a spike in quantity demanded (or may allow the buyer to lower *its* price enough to beat out *its* competitor, giving that buyer and consequently the supplier

[10]If labor were supplied elastically, the only differences between the planner's allocation and the decentralized equilibrium would be that less labor would be supplied because of the monopoly markups and all production would scale down in proportion to the decrease in aggregate labor.

a spike in quantity demanded). If two part tariffs are not available, solving for the optimal prices (and consequently the allocation) is difficult.

2 Size Distribution

Here I discuss two dimensions of the cross sectional distribution of size. First I give expressions for the conditional and unconditional distributions of number of customers (other firms that purchase intermediate goods). Second, I describe the cross sectional distribution of employment.

2.1 Number of Customers

Consider a single draw of a technique that uses good i as an input. Given the efficiency, q_i, we can compute the probability that the technique is the buyer's best available technique. To do this, we first characterize the following object: For a potential customer that has drawn a technique that uses i, what is the probability that it has no other techniques better than some efficiency level q?

Given the Poisson distribution over the number of techniques, a firm will have $n-1$ other techniques with probability $\frac{e^{-\tilde{\lambda}}\tilde{\lambda}^n}{n!}$. The CDF of efficiency delivered by each of these techniques if $G(q)$. We can therefore write the probability that the potential buyer has no other technique that delivers better than q as:

$$\sum_{n=1}^{\infty}\frac{e^{-\tilde{\lambda}}\tilde{\lambda}^n}{n!}G(q)^{n-1}=\frac{1}{G(q)}\left[\sum_{n=0}^{\infty}\frac{e^{-\tilde{\lambda}}\tilde{\lambda}^n}{n!}G(q)^n-e^{-\tilde{\lambda}}\right]=\frac{F(q)-e^{-\tilde{\lambda}}}{G(q)}$$

Among techniques that use i as a supplier, the fraction that deliver efficiency less than q is $H\left(\frac{q}{q_i^{\alpha}}\right)$, with density $\frac{1}{q_i^{\alpha}}H'\left(\frac{q}{q_i^{\alpha}}\right)$. We can now characterize the probability that a particular technique is the potential buyer's best technique:

$$\Pr\left(i \text{ is best}|q_i\right) = \int_0^\infty \frac{1}{q_i^\alpha} H'\left(\frac{\tilde{q}}{q_i^\alpha}\right) \frac{F(\tilde{q}) - e^{-\tilde{\lambda}}}{G(\tilde{q})} d\tilde{q}$$

How many known techniques use a given firm as a potential supplier? In other words, How many potential customers does a given supplier have? Across all firms, the distribution over the number of potential customers follows a Poisson law with with mean $\tilde{\lambda}$. Since each one of those techniques has an equal chance of being the potential buyer's best technique, so the distribution over the number of *actual* customers will also be a Poisson, with parameter:

$$\tilde{\lambda} \int_0^\infty \frac{1}{q_i^\alpha} H'\left(\frac{\tilde{q}}{q_i^\alpha}\right) \frac{F(\tilde{q}) - e^{-\tilde{\lambda}}}{G(\tilde{q})} d\tilde{q}$$

Using the functional form for H and taking the limit as $z_0 \to 0$ yields

$$\lambda \frac{q_i^{\alpha\zeta}}{\theta}$$

So among firms with efficiency q, the distribution over the number of customers will be given by a Poisson distribution with parameter $\lambda \frac{q^{\alpha\zeta}}{\theta}$. One can see that the distribution among high efficiency suppliers first order stochastically dominates the distribution among low efficiency suppliers: high efficiency firms get more customers.

We next look at the unconditional distribution over the number of customers among all firms.

To find the mass of suppliers with n customers, we integrate over suppliers of each efficiency:

$$\begin{aligned}
\int_0^\infty \frac{\left(\frac{\lambda}{\theta}q^{\alpha\zeta}\right)^n e^{-\frac{\lambda}{\theta}q^{\alpha\zeta}}}{n!}dF\left(q\right) &= \int_0^\infty \frac{\left(\frac{\lambda\theta^\alpha}{\theta}\left(\theta q^{-\zeta}\right)^{-\alpha}\right)^n e^{-\left(\frac{\lambda\theta^\alpha}{\theta}\left(\theta q^{-\zeta}\right)^{-\alpha}\right)}}{n!}\zeta\theta q^{-\zeta-1}e^{-\theta q^{-\zeta}}dq \\
&= \int_0^\infty \frac{\left(\frac{w^{-\alpha}}{\Gamma(1-\alpha)}\right)^n e^{-\left(\frac{w^{-\alpha}}{\Gamma(1-\alpha)}\right)}}{n!}e^{-w}dw
\end{aligned}$$

The second line uses the fact that $\theta = \Gamma(1-\alpha)\lambda\theta^\alpha$.

It is notable that this distribution depends on only one parameter, α (the share of inputs in production). This is because buyers choose to use the technique that gives the best combination of efficiency (z) and input cost (q). α determines the relative importance of these two factors. Recall that the efficiency associated with a single technique is zq^α. If α is large, than the the share of inputs is higher, and the cost of inputs becomes relatively more important. An increase in α makes the techniques using high efficiency suppliers even more cost effective. In contrast, when α is low, more weight is put on the idiosyncratic productivity associated with the technique. Because the productivity draws are drawn from the same distribution regardless of the efficiency of the supplier, lowering α increases the odds that a low efficiency firm will be able to attract customers; the low efficiency becomes less relevant to its customers.

Figure 1 shows the distribution of customers for different values of α. When α is high, more weight is put on the cost of inputs, so the distribution is more skewed. The tail is thicker, and but there are also more firms without any customers. In contrast, when α is low, the middle of the distribution is thicker.

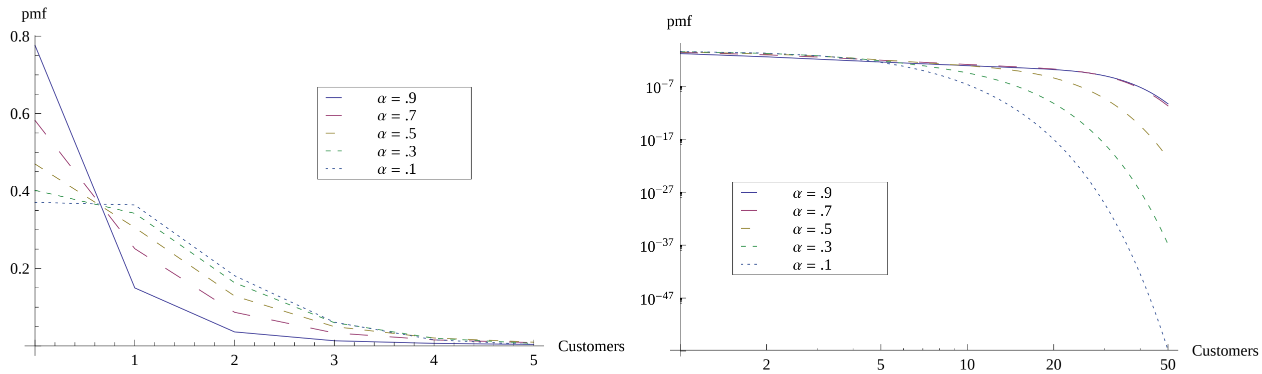

(a) Number of Customers

(b) Number of Customers, Log-Log

Figure 1: The Distribution of Customers

Figure 1a gives the mass of firms with n customers for several different values of the input share α. Figure 1b also gives the mass of firms with n customers, but on a log-log plot, to better show the tail of the distribution.

2.2 Distribution of Employment

In characterizing the cross sectional distribution of employment, we first derive a convenient fact. Let $B\left(q|q_i\right)$ be the the CDF of the efficiency of customers of suppliers with efficiency q_i. In other words, $B\left(q|q_i\right)$ is the distribution efficiency among firms whose best technique uses a supplier with efficiency q_i. We will show that in the limit as $z_0 \to 0$, $B\left(q|q_i\right) = F(q)$. We can solve for this distribution of customers' actual efficiency with an application of Bayes rule:

$$B'\left(q|q_i\right) \quad = \quad \frac{\frac{1}{q_i^\alpha} H'\left(\frac{q}{q_i^\alpha}\right) \frac{F(q)-e^{-\tilde{\lambda}}}{G(q)}}{\int_0^\infty \frac{1}{q_i^\alpha} H'\left(\frac{\tilde{q}}{q_i^\alpha}\right) \frac{F(\tilde{q})-e^{-\tilde{\lambda}}}{G(\tilde{q})} d\tilde{q}}$$

The numerator is the density of efficiency delivered by techniques that use i, $\frac{1}{q_i^\alpha} H'\left(\frac{q}{q_i^\alpha}\right)$, multiplied by the probability that a such a technique is the potential customer's best technique. The denominator is the probability that the the a technique from i is the customer's best technique. We can use the functional forms for H and take a limit as $z_0 \to 0$:

$$\lim_{z_0 \to 0} B'(q|q_i) = \lim_{z_0 \to 0} \frac{\zeta z_0^{\zeta} q_i^{\alpha\zeta} q^{-\zeta-1} e^{-\theta q^{-\zeta}}}{\int_{z_0 q_i^{\alpha}}^{\infty} \zeta z_0^{\zeta} q_i^{\alpha\zeta} \tilde{q}^{-\zeta-1} e^{-\theta \tilde{q}^{-\zeta}} d\tilde{q}}$$

$$\zeta\theta q^{-\zeta-1} e^{-\theta q^{-\zeta}}$$

$$= \lim_{z_0 \to 0} F'(q)$$

Knowing the efficiency of a supplier gives no information about the identity of its customer. That is, there are no systematic differences between the customers of low and high efficiency suppliers.

This has two implications. First, it gives insight into the determinants of size. High efficiency firms will (on average) be larger because they have more customers, not because their customers are on average bigger.

Second we can treat the characteristics of customers as independent, identically distributed random variables. This will be helpful in several ways. Of particular use here is the fact that we can treat the size of a customer as an IID random variable.

Distribution of Employment

We are interested in adding together the labor used to make goods for final consumption and for intermediate use for each customer. Since the latter can be treated as random variables, it will be easiest to work with characteristic functions of the relevant distributions. To this end we define several objects.

First, let $\chi(s)$ be the characteristic function associated with the cross sectional distribution of employment. This is the central object of interest, but to get at it, we take several intermediate steps. Next, let $\chi^0(s|q)$ be the characteristic function for labor used for final demand. If $L^0(q)$ is the (deterministic) quantity of labor used for final demand, we have

$$\begin{aligned}\chi^0(s|q) &= \int_{-\infty}^{\infty} e^{isl}\delta\left(l - L^0(q)\right) dl = e^{isL^0(q)} \int_{-\infty}^{\infty} e^{isx}\delta(x)dx \\ &= e^{isL^0(q)}\end{aligned}$$

where δ is the Dirac delta function.

We showed above that the quantity of labor used by a single customer is an IID random variable. Recall also the convenient fact that if firm j uses L_j units of labor, j's supplier will use $L_{ij} = \alpha L_j$ units of labor to make the inputs for j. If L_j is an IID random variable, αL_j is as well. With this in mind, let $\chi^1(s)$ be the characteristic function associated with the labor required to make the inputs for a single customer.

$$\begin{aligned}\chi^1(s) &= \int_{-\infty}^{\infty} \Pr\left(L_{ij} = l\right) e^{-isl} dl = \int_{-\infty}^{\infty} \frac{1}{\alpha} \Pr\left(L_j = \frac{l}{\alpha}\right) e^{-isl} dl \\ &= \int_{-\infty}^{\infty} \Pr\left(L_j = \frac{l}{\alpha}\right) e^{-i(\alpha s)\frac{l}{\alpha}} d\left(\frac{l}{\alpha}\right) \\ &= \chi(\alpha s)\end{aligned}$$

Let $\chi^{int}(s|q)$ be the characteristic function associated with the labor used for all intermediates among firms with efficiency q. Using the fact that the characteristic function of the sum of independent random variables is the product of the characteristic functions of each of the random variables, we can write $\chi^{int}(s|q)$ as

$$\begin{aligned}\chi^{int}(s|q) &= \sum_{n=0}^{\infty}\chi^1(s)^n \Pr(n \text{ customers}|q) = \sum_{n=0}^{\infty}\chi^1(s)^n \frac{\left(\frac{\lambda}{\theta}q^{\alpha\zeta}\right)^n e^{-\lambda\frac{q^{\alpha\zeta}}{\theta}}}{n!} = e^{-\Lambda\frac{q^{\alpha\zeta}}{\theta}\left[1-\chi^1(s)\right]} \\ &= e^{-\Lambda\frac{q^{\alpha\zeta}}{\theta}[1-\chi(\alpha s)]}\end{aligned}$$

where the second equality uses the fact that among firms with efficiency q the distribution of customers is Poisson with parameter $\lambda\frac{q^{\alpha\zeta}}{\theta}$.

We can put these together to derive an expression for $\chi(s|q)$, the characteristic function associated with the distribution of employment among firms with efficiency q:

$$\chi(s|q) = \chi^0(s|q)\,\chi^{int}(s|q) = e^{isL^0(q)}e^{-\frac{\lambda}{\theta}q^{\alpha\zeta}[1-\chi(\alpha s)]}$$

Lastly we can integrate across firms, which delivers a single recursive equation that defines $\chi(s)$:

$$\chi(s) = \int_0^\infty \chi(s|q)\,dF(q) = \int_0^\infty e^{isL^0(q)}e^{-\frac{\lambda}{\theta}q^{\alpha\zeta}[1-\chi(\alpha s)]}dF(q)$$

We now plug in the functional forms $L^0(q) = \frac{\left(\theta q^{-\zeta}\right)^{-\frac{\varepsilon-1}{\zeta}}}{\Gamma\left(1-\frac{\varepsilon-1}{\zeta}\right)}(1-\alpha)L$ and $dF(q) = \theta\zeta q^{-\zeta-1}e^{-\theta q^{-\zeta}}dq$ to give

$$\chi(s) = \int_0^\infty e^{is\left(\frac{\left(\theta q^{-\zeta}\right)^{-\frac{\varepsilon-1}{\zeta}}}{\Gamma\left(1-\frac{\varepsilon-1}{\zeta}\right)}(1-\alpha)L\right)}e^{-\Lambda\frac{\theta^\alpha\left(\theta q^{-\zeta}\right)^{-\alpha}}{\theta}[1-\chi(\alpha s)]}\theta\zeta q^{-\zeta-1}e^{-\theta q^{-\zeta}}dq$$

and using a change of variables along with the fact that $\theta = \Gamma(1-\alpha)\lambda\theta^\alpha$ gives

$$\chi(s) = \int_0^\infty \exp\left\{is\frac{t^{-\frac{\varepsilon-1}{\zeta}}}{\Gamma\left(1-\frac{\varepsilon-1}{\zeta}\right)}(1-\alpha)L - \frac{t^{-\alpha}}{\Gamma(1-\alpha)}[1-\chi(\alpha s)] - t\right\}dt \tag{14}$$

There are several things to note. First, as L changes, the distribution of labor scales uniformly.[11] Second the distribution depends on only two parameters, α and $\frac{\varepsilon-1}{\zeta}$. The share of intermediates matters for the same reason as before, it determines the skew of the distribution of number of customers. $\frac{\varepsilon-1}{\zeta}$ is a composite of two parameters, the elasticity of substitution in final consumption, and ζ, the tail index of both $H(\cdot)$ (the Pareto distribution from which productivity shocks are drawn) and $F(\cdot)$ (the cross sectional distribution of efficiency). In combination, these parameters determine the skewness of the distribution of final consumption. When ζ is small, the efficiency distribution has a thicker tail, inducing a thicker tail in the distribution of final consumption. When ε is high, consumers are more willing to substitute toward low cost goods, also thickening the tail of final consumption.

Equation (14) can be used to solve for $\chi(s)$ numerically.[12] We can consequently use standard methods to back out the distribution of employment form its characteristic function.

Figure 2 shows the distribution for the parameters $\alpha = \frac{\varepsilon-1}{\zeta} = 1/2$ and $L = 1$. One can see that this density is quite skewed, with the mode well below the mean of 1.

3 Asymmetric Networks

The previous analysis studied a very specific type of network, and leaves open the question of how alternative network configurations would affect aggregate productivity. For example, if many

[11]In fact, one could write the characteristic function associated with the fraction of labor used by each firm as $\hat{\chi}(s) = \chi\left(\frac{1}{L}s\right)$. This satisfies

$$\hat{\chi}(s) = \int_0^\infty \exp\left\{is\frac{t^{-\frac{\varepsilon-1}{\zeta}}}{\Gamma\left(1-\frac{\varepsilon-1}{\zeta}\right)}(1-\alpha) - \frac{t^{-\alpha}}{\Gamma(1-\alpha)}\left[1-\hat{\chi}(\alpha s)\right] - t\right\} dt$$

[12]While equation (14) is a functional equation, it resembles a difference equation. One can solve this using a reverse shooting algorithm, starting near the point $\chi(0) = 1$ and interpolating.

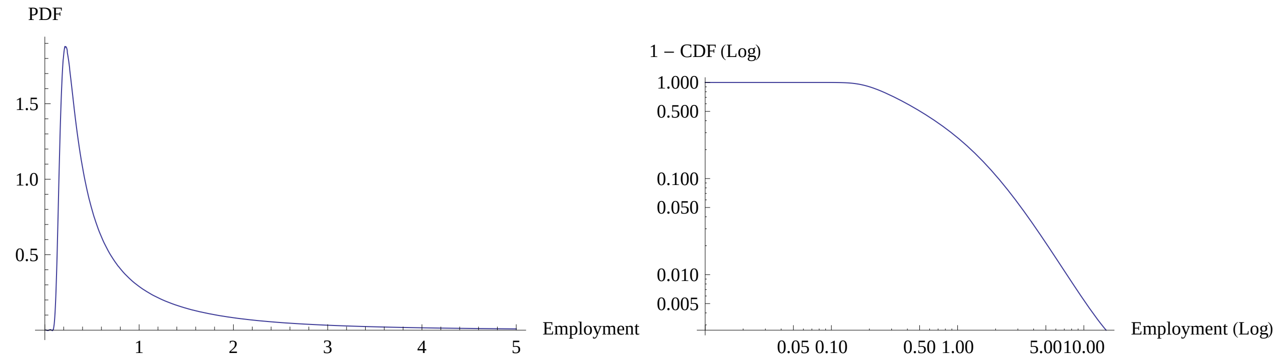

(a) Density of Employment

(b) Right CDF of Employment, Log-Log

Figure 2: The Distribution of Employment

Figure 2a gives the density of employment with $\alpha = \frac{\varepsilon-1}{\zeta} = 1/2$. Figure 2b gives right CDF on a log-log plot to show the shape of the tail of the distribution.

firms are in the center of a city and others are in the outskirts, how would aggregate productivity respond from this increased concentration? If some entrepreneurs were particularly active in finding new techniques (and others particularly inactive), how would this change patterns of diffusion of productivity gains from newly discovered techniques?

To get at these, we first extend the previous setup to allow for more interesting network configurations. There are multiple types of firms, indexed by $t \in T$. The only structural difference between firms of different types is how frequently they are involved with new techniques that are discovered. Let M_t be the mass of each type t firms and, abusing notation, let T be the number of types (in addition to the set of types).

The social planner's problem is exactly the same as in Section 1.2 and all first order conditions carry over. Instead of characterizing the distribution of efficiency across all firms, it will be convenient to characterize the distribution among each type. Let $F_t(q)$ be the fraction of type t firms with efficiency less than q. We proceed to characterize these distributions by setting up a fixed point problem.

We next define several objects that have analogs in Section 1. $\omega_t(n; t')$ is the fraction type t

firms that have access to n techniques with suppliers of type t'. At a given point in time this follows a Poisson distribution with mean $\tilde{\lambda}_t(t')$. Let $G_t(q;t')$ be the distribution of efficiency provided by a single technique drawn by a type t firm with supplier type t'. $G_t(q;t')$ is then:

$$G_t(q;t') = \int_0^\infty F_{t'}\left(\left(\frac{q}{z}\right)^{\frac{1}{\alpha}}\right) dH(z)$$

Given $G_t(q;t')$, we can compute $\Pr(Q_t < q)$, the probability that all draws from all types are less than some level q for a firm of type t:

$$\Pr(Q_t < q) = \prod_{t' \in T} \sum_{n=0}^{\infty} \omega_t(n;t') G_t(q;t')^n = e^{-\sum_{t'} \tilde{\lambda}_t(t')[1-G_t(q;t')]}$$

The same abuse of the law of large numbers gives $F_t(q) = \Pr(Q_t < q)$, giving the fixed point problem, T functional equations for the T unknown functions $\{F_t(\cdot)\}_{t \in T}$:

$$\log F_t(q) = -\sum_{t' \in T} \tilde{\lambda}_t(t') \left[1 - \int_0^\infty F_{t'}\left(\left(\frac{q}{z}\right)^{\frac{1}{\alpha}}\right) dH(z)\right] \tag{15}$$

3.1 A Parametric Assumption

We will use the same functional forms as in the one type model, $H(z) = 1 - \left(\frac{z}{z_0}\right)^{-\zeta}$ and $\tilde{\lambda}_t(t') = \lambda_t(t') z_0^{-\zeta}$. We will then look at the equilibrium of the limiting economy as $z_0 \to 0$. With a similar argument, we will show that any set of solutions to equation (15) will follow Frechet distributions. As before we can write

$$\tilde{\lambda}_t(t')\left[1 - G_t(q;t')\right] = q^{-\zeta} \lambda_t(t') \int_0^{\left(\frac{q}{z_0}\right)^{\frac{1}{\alpha}}} \alpha\zeta x^{\alpha\zeta-1}\left(1 - F_{t'}(x)\right) dx$$

Substituting into equation (15) for each t' we get

$$-\log F_t(q) = q^{-\zeta} \sum_{t' \in T} \lambda_t(t') \int_0^{\left(\frac{q}{z_0}\right)^{\frac{1}{\alpha}}} \alpha\zeta x^{\alpha\zeta - 1} \left(1 - F_{t'}(x)\right) dx$$

For any set of $\{F_t(\cdot)\}_{t \in T}$, as $z_0 \to 0$, this expression goes to $q^{-\zeta}$ multiplied by a constant. For each t, label this constant θ_t so that $F_t(q) = e^{-\theta_t q^{-\zeta}}$.

We next solve for $\{\theta_t\}_{t \in T}$, which are defined to satisfy

$$\theta_t = \sum_{t' \in T} \lambda_t(t') \int_0^{\infty} \alpha\zeta x^{\alpha\zeta - 1} \left(1 - F_{t'}(x)\right) dx = \sum_{t' \in T} \lambda_t(t') \int_0^{\infty} x^{\alpha\zeta} F'_{t'}(x) dx$$

Plugging in the functional form for $F_{t'}(q)$ gives

$$\theta_t = \Gamma(1 - \alpha) \sum_{t' \in T} \lambda_t(t') \theta_{t'}^{\alpha} \tag{16}$$

Notice also that, as before, for any $\{\lambda_t(t')\}_{t,t' \in T}$ there are three solutions to equation (16): $\theta_t = 0, \forall t \in T$, $\theta_t = \infty, \forall t \in T$, and a third solution that is the solution to the planner's problem.

3.2 Aggregate Output

Given the distribution of efficiency across firms, total output will be $Y^0 = QL$ where $Q^{\varepsilon-1} = \int_J q_j^{\varepsilon-1}$ (again, the analysis in the one type economy carries over). It will be convenient to define $Q_t \equiv \left(\int_0^{\infty} q^{\varepsilon-1} dF_t q\right)^{\frac{1}{\varepsilon-1}}$ to be a productivity aggregator among firms of type t. We can then write

$$Q^{\varepsilon-1} = \sum_t M_t Q_t^{\varepsilon-1}$$

With the functional forms, we can write (as before):

$$Q_t = \theta_t^{1/\zeta}\Gamma\left(1 - \frac{\varepsilon - 1}{\zeta}\right)^{\frac{1}{\varepsilon-1}}$$

The productivity aggregator for the whole economy can then be written as

$$Q^{\varepsilon-1} = \Gamma\left(1 - \frac{\varepsilon - 1}{\zeta}\right)\sum_{t\in T} M_t \theta_t^{\frac{\varepsilon-1}{\zeta}}$$

4 Superstars and Productivity Spillovers

The purpose of this section is to demonstrate how the configuration of the network affects aggregate output. We consider examples in which there are two types of firms, indexed by $t \in \{A, B\}$. The total mass of firms is $M_A + M_B = 1$. Again, the types of firms differ only in how frequently they are involved with new techniques that are discovered. For a firm of type t let λ_t be the (cumulative) arrival rate of techniques, so that $\lambda_t = \sum_{t'} \lambda_t(t')$.

To focus on the influence of the configuration, I will hold the total number of techniques in the network constant but vary the distribution of links. In other words, I will let $\{\lambda_t(t')\}$ and $\{M_t\}$ vary subject to

$$\lambda = M_A\left(\lambda_A(A) + \lambda_A(B)\right) + M_B\left(\lambda_B(A) + \lambda_B(B)\right)$$

holding λ constant .

The heterogeneity across types is parameterized as follows: Let $\rho_p \equiv \frac{\lambda_A}{\lambda_B} = \frac{\lambda_A(A)+\lambda_A(B)}{\lambda_B(A)+\lambda_B(B)}$ (p is for "producer"). This is a measure of how much more frequently type A firms discover new techniques. If $\rho_p > 1$ type A firms discover new techniques faster.

Similarly, let $\rho_s \equiv \frac{(\lambda_A(B)+\lambda_B(A))/M_A}{(\lambda_A(B)+\lambda_B(B))/M_B}$ (s is for supplier). This is a measure of how much more frequently type A firms are suppliers for new techniques that get discovered. If $\rho_s = 1$ then the probability of being a supplier is uniform across all firms. If $\rho_s > 1$, then a new technique is relatively more likely to use a type A firm than would be suggested by M_A and M_B.

Given the values of ρ_p and ρ_s (along with the assumption that $\frac{\lambda_A(B)}{\lambda_A(A)} = \frac{\lambda_B(B)}{\lambda_B(A)}$ which can easily be abandoned) we can solve for the implied values of $\{\lambda_A(A), \lambda_A(B), \lambda_B(A), \lambda_B(B)\}$:

$$\begin{aligned}
\lambda_A(A) &= \frac{\rho_s M_A}{\rho_s M_A + M_B} \frac{\rho_p}{\rho_p M_A + M_B} \lambda \\
\lambda_A(B) &= \frac{M_B}{\rho_s M_A + M_B} \frac{\rho_p}{\rho_p M_A + M_B} \lambda \\
\lambda_B(A) &= \frac{\rho_s M_A}{\rho_s M_A + M_B} \frac{1}{\rho_p M_A + M_B} \lambda \\
\lambda_B(B) &= \frac{M_A}{\rho_s M_A + M_B} \frac{1}{\rho_p M_A + M_B} \lambda
\end{aligned}$$

To compute aggregate output, we need an expression for the productivity aggregator $Q = \left(\int_J q_j^{1-\frac{1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}$. Given the values of $\{\lambda_t(t')\}_{t,t' \in T}$ we can solve for the values of θ_A and θ_B using equation (16). With this, we can then solve for the productivity aggregator Q:

$$Q^{\varepsilon-1} = M_A Q_A^{\varepsilon-1} + M_B Q_B^{\varepsilon-1} \propto M_A \theta_A^{\frac{\varepsilon-1}{\zeta}} + M_B \theta_B^{\frac{\varepsilon-1}{\zeta}}$$

where Q_A and Q_B are the productivity aggregators among firms of each type. We are interested in how Q varies as the structure of the network changes.

Hubs

We first consider the case in which type A firms are hubs: they are more likely to discover new techniques and they are relatively more likely to be the supplier when other firms get new techniques, as shown in Figure 3. Formally we set $\rho_p = \rho_s$ and vary this common number for various values of M_A. The graph shows the aggregate output relative to the uniform network ($\rho_p = \rho_s = 1$). Note that aggregate output rises (as $\rho_p, \rho_s \to \infty$, relative productivity slowly falls back to 1 for all three curves). Aggregate output is higher for the following reason: When $\rho_p > 1$, type A firms draw many techniques and are therefore more productive. While type B firms draw a smaller number of techniques, the techniques they do draw are likely to have type A firms as suppliers, meaning that the type B firms are likely to get a low marginal cost. When $\rho_p, \rho_s \to \infty$, the impact disappears, because all type B firms are essentially disconnected from the network. Type A firms are more productive (there are more techniques among them) but there are fewer firms, and consequently less gains from variety. These effects exactly offset, and aggregate output is the same as the uniform case.

Figure 3: Varying Both ρ_p and ρ_s

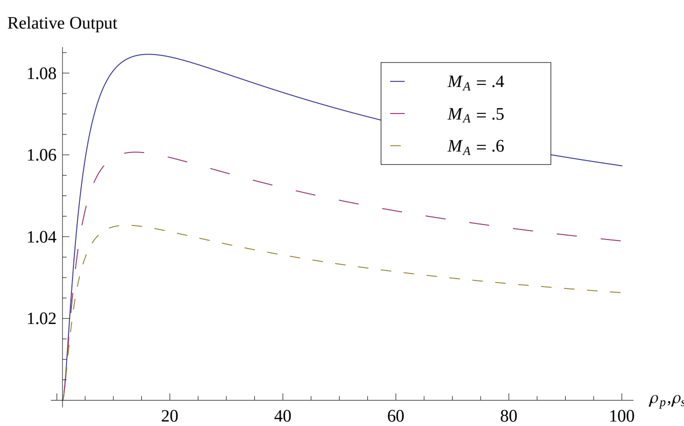

For intermediate values, the increase in aggregate output is larger when M_A is small. This happens because more techniques are concentrated within the type A firms, so productivity among those firms is high. Type B firms are likely to have techniques that use type A firms, so that they are increasingly able to benefit from the high productivity among type A that is due to the increased concentration. This is most stark as $M_A \to 0$, in which case the peak of the curve rises unboundedly. More formally, for a given M_A, set $\rho_p = \rho_s = \frac{1/2}{M_A}$. Then $\lim_{M_A \to 0} Q = \infty$.

An even more stark example (which is easier to analyze by hand) is $\rho_p = \frac{M_B}{M_A}$ and $\rho_s \to \infty$. This leads to (normalizing $\lambda = 1$), $\lambda_A(A) = \frac{1/2}{M_A}$, $\lambda_B(A) = \frac{1/2}{M_B}$, $\lambda_A(B) = \lambda_B(B) = 0$. Here, no matter how few type A firms there are, half of all total techniques are drawn by those firms. However, *all* techniques use type A firms as inputs. Again, we have $\lim_{M_A \to 0} Q = \infty$.

We next examine the case in which ρ_p varies but $\rho_s = 1$, shown in Figure 4. In this case type A firms are more likely to discover new techniques, but no firm is more likely to be the object of a new technique. Here we can see that total output is smaller than the uniform case for any configuration (and the drop is persistent: as $\rho_p \to \infty$, relative productivity stays depressed below 1). Here, while type A firms are more productive, as a group than they would be in the uniform case, the type B firms are less able to take advantage of this, because the techniques they find aren't especially concentrated on the type A firms. As a result, type B firms are less productive in than in the uniform case, so much so that this dominates the increased productivity among type A firms.

Figure 4: Varying ρ_p, $\rho_s = 1$

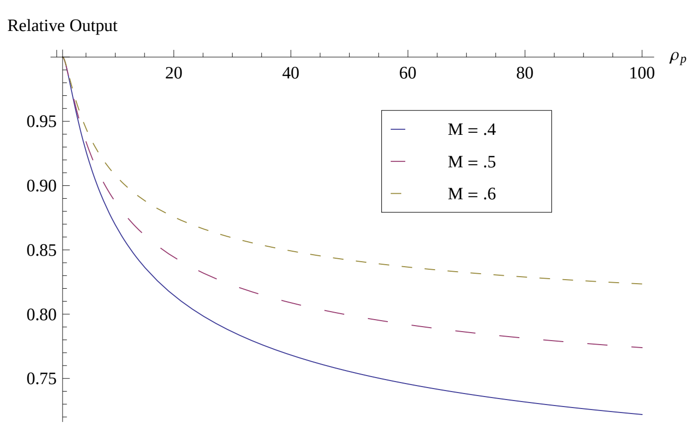

Lastly we examine the the case in which $\rho_p = 1$ but ρ_s varies, shown in Figure 5. Here productivity is exactly the same as in the uniform case. Considering this and the previous case, we can infer the different roles of ρ_p and ρ_s. ρ_p generates productivity differences across the different types, as drawing more techniques leads to (on average) higher productivity. ρ_s determines how much of these productivity differences spill over to the other types. In this case where $\rho_p = 1$, there are no productivity differences to spill over, so varying ρ_s makes no differences. Compare this to $\rho_p > 1$, in which case varying ρ_s can make a big difference.

Figure 5: $\rho_p = 1$, Varying ρ_s

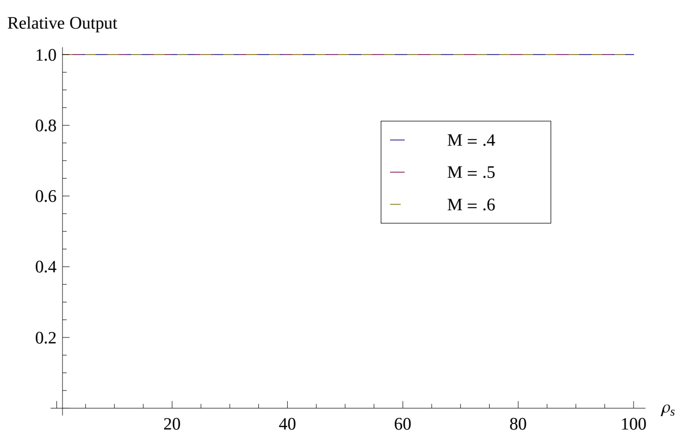

5 Conclusion

This paper has described a tractable model of the formation and evolution of chains of production. The model aggregates easily, with simple formulas connecting the density of the network to aggregate output. With more interesting network configurations, the model can be solved almost as easily. We find that, holding the number of techniques constant, the particular configuration can have an enormous impact on aggregate productivity.

References

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Cascades in Networks and Aggregate Volatility," NBER Working Papers 16516, National Bureau of Economic Research, Inc November 2010.

Alvarez, Fernando E., Francisco J. Buera, and Jr. Robert E. Lucas, "Models of Idea Flows," NBER Working Papers 14135, National Bureau of Economic Research, Inc June 2008.

B, Jr Long John and Charles I Plosser, "Real Business Cycles," *Journal of Political Economy*, February 1983, *91* (1), 39–69.

Carvalho, Vasco, "Aggregate fluctuations and the network structure of intersectoral trade," Economics Working Papers 1206, Department of Economics and Business, Universitat Pompeu Fabra November 2007.

Dupor, Bill, "Aggregation and irrelevance in multi-sector models," *Journal of Monetary Economics*, April 1999, *43* (2), 391–409.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, September 2002, *70* (5), 1741–1779.

Foerster, Andrew T., Pierre-Daniel G. Sarte, and Mark W. Watson, "Sectoral vs. Aggregate Shocks: A Structural Factor Analysis of Industrial Production," NBER Working Papers 14389, National Bureau of Economic Research, Inc October 2008.

Horvath, Michael, "Cyclicality and Sectoral Linkages: Aggregate Fluctuations from Independent Sectoral Shocks," *Review of Economic Dynamics*, October 1998, *1* (4), 781–808.

Jones, Charles I., "Intermediate Goods, Weak Links, and Superstars: A Theory of Economic Development," NBER Working Papers 13834, National Bureau of Economic Research, Inc March 2008.

Kelly, Morgan, "The Dynamics of Smithian Growth," *The Quarterly Journal of Economics*, August 1997, *112* (3), 939–64.

_ , "The Invention of Invention," Working Papers 200515, School Of Economics, University College Dublin September 2005.

Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, November 1997, *65* (6), 1389–1420.

Lucas, Robert E., "Ideas and Growth," *Economica*, 02 2009, *76* (301), 1–19.