Input Diffusion and the Evolution of Production Networks*

Vasco Carvalho
Cambridge and CREI

Nico Voigtländer
UCLA and NBER

First Version: March 2014

Abstract

What determines which inputs are initially considered and eventually adopted in the production of new or improved goods? Why are some inputs much more prominent than others? We model the evolution of input linkages as a process where new producers first search for potentially useful inputs and then decide which ones to adopt. A new product initially draws a set of 'essential suppliers'. The search stage is then confined to the network neighborhood of the latter, i.e., to the inputs used by the essential suppliers. The adoption decision is driven by a tradeoff between the benefits accruing from input variety and the costs of input adoption. This has important implications for the number of forward linkages that a product (input variety) develops over time. Input diffusion is fostered by network centrality – an input that is initially represented in many network neighborhoods is subsequently more likely to be adopted. This mechanism also delivers a power law distribution of forward linkages. Our predictions continue to hold when varieties are aggregated into sectors. We can thus test them, using detailed sectoral US input-output tables. We show that initial network proximity of a sector in 1967 significantly increases the likelihood of adoption throughout the subsequent four decades. The same is true for rapid productivity growth in an input-producing sector. Our empirical results highlight two conditions for new products to become central nodes: initial network proximity to prospective adopters, and technological progress that reduces their relative price. Semiconductors met both conditions.

JEL: O33, C67, D57, L23

Keywords: Input adoption, directed network search, dynamics of input-output networks

*We would like to thank Francesco Caselli, Jeremy Greenwood, Ezra Oberfield, Giacomo Ponzetto, and Jaume Ventura, for helpful comments and suggestions. Voigtländer acknowledges support from the UCLA Anderson Easton Technology Leadership Program.

1 Introduction

Technology is typically formalized as a 'production recipe', prescribing a combination of different inputs to produce output. For a given individual producer, this reflects a set of supplier relationships. Across all products, these buyer-supplier relationships give rise to the input-output structure of the economy. Consequently, each individual producer is embedded in a network of production linkages – it does not only interact directly with its suppliers, but also indirectly with input producers further upstream. By the same token, technological progress is associated with the evolution of 'production recipes', and the adoption of new inputs constitutes a defining characteristic of this process. What determines which inputs are initially considered and eventually adopted in the production of new or improved goods? Why are some inputs much more prominent than others?

In this paper we study input adoption through a network perspective. This approach gives rise to the notion of network distance between any *potential* buyer-supplier pair, i.e., a producer's distance to inputs that are not (yet) included in its 'production recipe'. We show that this novel view has important implications for the adoption of new inputs: individual producers are more likely to adopt inputs that are closer in their network neighborhood. This is similar to the evolution of social networks, where new friendships are more likely to form with friends' friends than with random people. In the production context, a firm is more likely to adopt an input that its supplier is already using, than a random input from the product universe. Our results thus imply that the current production network and its evolution are closely interrelated: On the one hand, present network distances determine input adoption; input adoption, in turn, changes network distances. This gives rise to a dynamic evolution of the input-output network.

We begin by building a model of network formation at the variety level, where producers search for potentially useful inputs within their network neighborhood. The model predicts that initially closer network proximity implies higher likelihood of input adoption. Additionally, the model delivers a power law in the number of varieties supplied. We show that both these predictions continue to hold when varieties are aggregated into sectors based on input similarity – the standard rule for sectoral classification. The sector-level aggregation allows us to test these predictions in US input-output tables. We show that the initial network structure in 1967 predicts the formation of new linkages in the following four decades. In particular, sectors are more likely to adopt inputs to which they are already more closely (but indirectly) connected via their existent suppliers.

Our variety level model is motivated by the fact that input-output tables ultimately reflect transactions between individual producers. We represent the network by taking product varieties as nodes and input purchases as directed edges. Every period, a new variety emerges exogenously. It then forms input linkages following three steps, where the first two build on the central mecha-

nism of dynamic network formation models (Vázquez, 2003; Jackson and Rogers, 2007; Chaney, 2013). First, a new variety draws a set of essential input suppliers (or 'network parents') at random. Second, the new variety producer identifies further potential inputs by a local search mechanism, following the linkages of its essential suppliers. In other words, the search is directed towards the technological neighborhood of essential inputs. Third, the new variety producer decides which inputs to adopt among those identified in the second step. This decision is driven by a trade-off between benefits from a larger set of input varieties (à la Romer, 1990) and variety-specific customization costs for each adopted input. As a result, a finite optimal number of inputs is adopted from the network neighborhood of essential inputs. Products with more pre-existing forward linkages are more likely to be in the neighborhood of any (randomly drawn) essential inputs. Thus, they have a higher probability of being adopted by a new variety. In this way, already central nodes tend to build more forward linkages and thus become ever more central. Similar to the social network literature, our directed search mechanism gives rise to superstar varieties as reflected by a power law distribution.

We then explore the sector-level implications of such a mechanism. To define sectors in the model, we build on the rules by which new commodities are assigned to sectors in actual input-output tables. This classification is based on a variety's essential inputs. For example, a new variety that draws tires, an engine, and a body as essential inputs will be assigned to the motor vehicles sector. We show that, based on this definition, the model predicts i) the power law distribution of forward linkages aggregates up to the sector-level and ii) new input linkages across sectors are more likely to emerge within the proximity of existent input supply relations. Thus, even if the underlying network formation is happening at the *variety* level, we can make use of *sectoral* input-output data to examine the mechanism at work.

To test the predictions of our model, we use U.S. input-output tables at the 4-digit level between 1967 and 2002. Based on the observed intersectoral linkages in manufacturing, we compute a standard measure of network distance between any sector pair.[1] We find that sectors are substantially more likely to adopt inputs that are initially closer in their input-output network. This finding holds both in a panel setting where the input-output network evolves over time, and also in a cross-sectional analysis showing that closer network proximity in 1967 reduces the subsequent time to adoption. Our results are robust to a host of controls such as fixed effects for adopting and

[1]Technically, network distance is the shortest path from a producer j to input i. Practically, this can reflect several dimensions. First, technological distance in the sense that production processes are more or less similar. For example, engines are technologically closer to vehicles than processed food. Second, spatial distance to the extent that industries that trade inputs intensively tend to coagglomerate, and third, it can reflect idea flows and R&D spillovers (Ellison, Glaeser, and Kerr, 2010). For example, sectors that trade more intensively also have more intensive cross-citation patterns.

input-producing sectors. They are also economically significant: a one-standard deviation (std) decrease in network distance raises the adoption probability in any given benchmark year (5-year periods) by one third. In addition, we show that rapid technological progress in input-producing sectors makes their adoption more likely. Because technological progress may be driven by adoption, we use total factor productivity (TFP) growth prior to our sample period (1958-67) to predict subsequent productivity growth. These results confirm our findings.

We build on a rich research agenda that has studied input adoption and diffusion. Starting from the seminal work by Griliches (1957), a large and diverse literature has studied the diffusion of technology. A macro strand of this literature has focused on how particular technologies – such as electricity or semi-conductors – are progressively adopted by an expanding range of sectors. This gives rise to General Purpose Technologies (GPT) that mark historical eras and are seen as engines of growth (Helpman and Trajtenberg, 1998; Jovanovic and Rousseau, 2005). As in this literature, we are interested in understanding how a particular technology can emerge as an input supplier to many other technologies.[2] Our results imply two factors that raise the odds for new technologies to become GPTs: first, a relatively central position in the network. This condition is met if the new input is used by other prominent technologies. For example, in 1967 semiconductors were used as an input to electronic components – a prominent technology that in turn was used by a large number of other sectors. The second condition is rapid technological progress, so that the price of the new input falls, which makes its adoption more attractive. Intuitively, a central position in the network helps a new input to be 'visible' to potential adopters, and falling prices make eventual adoption more likely.

Our paper is also related to a micro strand of the literature that focuses on the role of social networks in the adoption of particular technologies (c.f. Young, 2003; Conley and Udry, 2010; Banerjee, Chandrasekhar, Duflo, and Jackson, 2013). We share the view that the adoption of technologies is mediated through a network. However, rather than focusing on the role of local social interactions, we study the importance of distance in the technological network more broadly.[3]

Our focus on input-output networks is also motivated by an emerging literature (Carvalho, 2010; Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi, 2012; Bigio and La'O, 2013) emphasiz-

[2]Interestingly, while Helpman and Trajtenberg (1998) rationalize the staggered diffusion of a GPT in terms of asymmetric adoption costs, they also conjecture that the order of adoption could be the result of "linkages between adopting sectors" and thus, that "technological proximity" may be an important factor in explaining diffusion patterns of GPTs. Our key mechanism formalizes this notion of "technological proximity" by placing technologies in a network and emphasizing network proximity as a key driver of adoption.

[3]Of course, distance in the input-output network may also reflect less frequent social interaction. For example, a tire producer is more likely to interact with people from the automotive industry than with pharmaceutical staff. Our argument exploits the variation across *sector pairs*, whereas the micro literature on social networks examines the role of local social interactions for the adoption of a *given* technology.

ing the role of the network structure of intersectoral linkages in propagating idiosyncratic shocks throughout the economy and generating aggregate fluctuations.[4] This literature invariably takes the input-output network as an exogenously given restriction on sectoral production technologies – or who sources inputs from whom – in order to study the impact of network structure on the strength of the shock propagation mechanism. Finally, our work builds on a literature of dynamic network formation models (Jackson and Rogers, 2007; Chaney, 2013). As in these papers, our network evolution process stresses the fact that existing links can be used to find new links: goods producers probe their existing set of input suppliers to find other potentially useful varieties for their own production process. In this context, our paper is also closely related to Oberfield (2012). While Oberfield also studies the formation of production networks over time, he does not exploit the underlying network structure to explain link formation. Instead, his mechanism has producers *randomly* searching for the lowest cost input supplier, while we emphasize the role of networks in the search for potential inputs.

Relative to the existing literature we make several contributions. First, relative to the literature on input adoption and input diffusion, we are the first to exploit the role of the input-output network structure in shaping the future path of input adoption, both theoretically and empirically. Second, relative to the literature on input-output networks and aggregate fluctuations, we endogenize the formation of input-output linkages. We thus provide a deeper understanding of asymmetries across input suppliers – a crucial precondition for the emergence of aggregate fluctuations. Third we show that the network search mechanism from the literature on dynamic network formation can also shed light on processes of input adoption and the evolution of technology. Fourth, we provide strong empirical support for our network based mechanism of input adoption. Finally, our theoretical and empirical findings have important implications for the rise of General Purpose Technologies – we show that both network centrality and rapid technological progress are necessary conditions in this process.

The paper is organized as follows. Section 2 uses the diffusion of semiconductors as a case study to illustrate our mechanism. Section 3 describes our model of input adoption, starting at product variety level and then aggregating these into sectors. Section 4 introduces our measure of network distance and describes our data. In Section 5 we present empirical results lending strong support to the predictions of our model. Section 7 concludes.

++[for related lit++] The process of link formation is not only a defining element of technological progress; it also determines the network structure of an economy. A growing literature stresses

[4]This work in turn builds on an older literature that emphasizes the role of input-output linkages for co-movement across sectors (c.f. Long and Plosser, 1983; Horvath, 1998; Conley and Dupor, 2003). See also Foerster, Sarte, and Watson (2011) and di Giovanni, Levchenko, and Mejean (2012) for empirical evidence supporting these mechanisms.

the importance of input-output linkages for macroeconomic outcomes. By propagating shocks, prominently linked sectors can create aggregate fluctuations (Acemoglu et al., 2012). Intersectoral linkages can also amplify idiosyncratic sectoral distortions into large aggregate productivity differences (Ciccone, 2002; Jones, 2013). It is therefore key to understand the evolution of the input-output structure and, in particular, why some sectors play a disproportionate role as input providers.

++ By this notion, there is a wealth of 'production recipes' that vary substantially across sectors – and even within sectors: firms can pick from a large pool of potentially suitable inputs to produce products within narrowly defined categories.[5]

++Correspondingly, changes in the input-output structure – input adoption – are at the heart of technological progress. What determines which inputs are initially considered and eventually adopted in the production of new or improved goods? Why are some inputs much more prominent than others?[6]

2 The Diffusion of Semiconductors

The diffusion of semiconductors, a key general purpose input, provides a telling illustration of input adoption in a network. Figure 1 shows a network representation of the US input-output table in 1967. Each 4-digit SIC sector is represented by a node, and edges between these nodes depict input flows across sectors. The solid black node on the left hand side of the graph corresponds to semi-conductors. The red nodes mark sectors that directly sourced semiconductors as an input in 1967 – only a handful of technologies incorporated semiconductors. Finally, the red arrows point to indirect users of semiconductors, i.e., sectors that sourced inputs which in turn used semiconductors.

Given this starting point, Figures 2-4 show the path of diffusion of semiconductors across sectors over the subsequent 15 years.[7] Blue dots in Figure 2 represent sectors that adopted semiconductors in 1972, as per the detailed input output tables of that year. Note that the new adopters also add new indirect paths to semi-conductors, as indicated by the blue lines in Figure 2. Cyan and green dots in Figures 3 and 4 correspond to sectors that adopted semiconductors by 1977 and 1982, respectively. As before, lines in the respective color represent newly formed indirect links. We ask whether these indirect linkages to semi-conductors are informative about the likelihood of

[5] Steve Jobs famously had the first iPhone's screen changed from plastic to hardened glass only four weeks before mass production began in 2007.

[6] The number of sectors that source inputs from a given supplier follows a power law (Carvalho, 2010). Atalay, Hortaçsu, Roberts, and Syverson (2011) and Kelly, Lustig, and Van Nieuwerburgh (2013) report evidence on the distribution of supply linkages at the firm level.

[7] Note that throughout we hold the 1967 network fixed. That is, all colored edges refer to input linkages observed in 1967.

subsequent direct adoption of semiconductors as an input.

[Insert Figures 1-4 here]

The pattern emerging from these Figures is striking. Every single one of the seven adopters in 1972 previously had an indirect connection to semiconductors via one other intermediate input. In the terminology of networks, all second-round adopters of semiconductors were two edges away (i.e., distance 2) from semiconductors. Similarly, four out of the five sectors that adopted semiconductors by 1977 sourced inputs from either the 1972 or the 1967 adopters. By 1982, the number of sectors using semiconductors as an input had more than trebled relative to 1967, setting the stage for the generalized adoption that would ensue in the 1990s and 2000s. Summarizing, the large majority (14 out of 17) of newly adopting sectors in this early 15 year period of diffusion of semiconductors were either distance 2 or 3 from semiconductors in 1967. In contrast, the typical 1967 network distance from semiconductors to a randomly selected sector in the economy was 5. That is, early adoption of semiconductors was strongly correlated with being in the network proximity of semiconductors early on.

It is instructive to focus on one of these paths of adoption to better understand the role on linkages across sectors in the diffusion of semiconductors. One of the earliest uses of semiconductors was in the invention and production of integrated circuits or chips, classified into the Electronic Components sector, one of our 1967 nodes in the network. While early computers sourced transistors from the Electronic Components sectors, the Computer and Office Equipment sector did not adopt semiconductors until new computer varieties made use of the newly developed integrated circuits. The world's first personal computer – the 'Kenbak-1' produced in 1970 – was the first computer device to source integrated circuits as an input (from the Electronic Components sector) and, alongside it, semiconductors. The Computers and Office Equipment sector thus appears in our network as one of the 1972 second round adopting sectors. Downstream of the Computers and Office Equipment sector we find the Scales and Balances sector, an adopter of semiconductors in 1977. This sector sourced early computer varieties in the late 1960s to store and perform calculations on weighing measurements. Throughout the 1970s the introduction of newer, smaller computer equipment varieties – itself made possible by the adoption of integrated circuits – opened the way for the large scale production of industrial and retail digital scales which themselves incorporated semiconductors directly as an input.

3 A model of input diffusion in a network of technologies

In this section we present a simple model of dynamic input diffusion across a network of interconnected product varieties. New varieties emerge exogenously every time period. Interconnections

across varieties reflect input needs, i.e., each variety is produced by incorporating other, already existent, varieties as intermediate inputs. These input linkages across varieties give rise to a network that evolves over time, as new varieties are introduced and new links are formed.

Building on the dynamic network formation models of Jackson and Rogers (2007) and Chaney (2013), we begin by modeling how the set of feasible inputs available to each new variety is defined. Following this literature, our network evolution process stresses the fact that existing links can be used to find new links. In our context, this means that a new variety is first assigned a set of 'essential' inputs and can then probe the network neighborhood of this set to find other varieties that can be of potential use as inputs.

Given this set of potential inputs available to each new variety, we proceed to endogenize the input adoption decision. We assume that input adoption is costly. Specifically, in order for a new variety to adopt an input, it must be customized at a cost that is specific to each variety-input pair. In the model, new variety producers face a trade-off between this customization cost and a love of variety effect accruing to adopting additional inputs. The solution to this tradeoff determines the total number of inputs that each new variety adopts.

Finally, in order to derive testable predictions that can be taken to sectoral input-output data, we explore the sector level implications of the variety level model. We classify varieties into sectors based on a principle of similarity of inputs that is also used in the construction of input-output tables. As a result, sectors are composed of varieties that share similar production processes, i.e., varieties that process similar input bundles. Based on this definition we can show that the key variety level mechanism – a new variety is more likely to adopt inputs in its network neighborhood – is still present after aggregation to the sectoral level. That is, among all pairs of sectors that are only indirectly linked to a given sector at a given point in time, those that are closer (in a network proximity sense) are more likely to be adopted as input providers subsequently.

3.1 Variety Level Model

Given a finite number of product varieties, t, we define a variety-level input-output matrix as a weighted directed network, represented by a $t \times t$ matrix where each entry $v_{ij} \geq 0$ denotes the flow of input variety i into variety j's production process. We say that j uses input i if $v_{ij} > 0$. Correspondingly, we define the unweighted directed network as the binary $t \times t$ matrix where each entry $b_{ij} \in \{0, 1\}$ denotes whether product variety j uses input variety i. To characterize the evolution of the variety-level network, we focus on b_{ij}, i.e., the formation of links.[8]

This production network evolves over time as new varieties arrive sequentially in the economy.

[8] Below, we show that under price symmetry, $v_{ij} \geq 0$ is proportional to b_{ij}.

In particular, at each time t a new variety is added to the economy.[9] Each new product variety initially draws a finite set K_t of necessary or 'essential' inputs; let m_K denote the number of input varieties in this set (for simplicity ignoring the subscript t). These draws occur uniformly at random across all existing varieties. Essential inputs can be thought of as defining features of the new variety. For example, if t is a car its set K_t will include a body, an engine, wheels, etc. There can be different varieties (or versions) of each essential input, but not all are necessarily used. In our example, the car producer may consider several different engine options.

In a second step, the new variety can adopt further inputs from the neighborhood of its essential suppliers. This reflects a stage of refinement of variety t by adding features beyond the essential ones. In the car example, the producer may look for options to make the body lighter or add luxury features to the car. We assume that the corresponding search for inputs is directed to the network neighborhood (input suppliers) of existing car body producers. For example, the BMW i3 has an ultra-light carbon fiber body – a material that has previously been used in Formula 1 cars. Similar to the first step, the second round search may also deliver a spectrum of potential inputs, and only a subset of these will eventually be adopted. Thus, the network search can be viewed as the producer of a new variety searching for which inputs (and thus production techniques) are technologically feasible.

To formalize the process of input search in the supplier network, let N_t denote the set of input varieties that producer t identifies as useful from its network search. This search follows the links of t's essential input suppliers in the set K_t. The number of varieties in the set N_t is denoted by m_N. One interpretation of this setup is that the network neighborhood of essential inputs defines which further varieties are *technologically* close to t and can therefore be of potential use in its production process. Alternatively, the setup can be interpreted as a local search process by which the developers of the new variety *learn* about other useful technologies via the personal interaction with their essential input suppliers.

We use this setup to study the probability with which a new variety t adopts a given input i. In the theory of network formation, this is related to the evolution of the outdegree of variety i.[10] The outdegree of each variety, $d_i^{out}(t)$, is heterogeneous across i and over time t. For an existing variety with outdegree $d_i^{out}(t)$ at time t, the expected growth rate of its outdegree is given by:

$$\frac{\partial d_i^{out}(t)}{\partial t} = p_K \frac{m_K}{t} + p_N \frac{m_K d_i^{out}(t)}{t} \frac{m_N}{m_K(p_K m_K + p_N m_N)} \tag{1}$$

[9] We use the index t to denote the new variety in each respective period. Thus, the index t refers to both the latest new variety that has been introduced, and the time period when this happened. Since varieties can be both inputs and output in our model, we use the notation 'input varieties' vs. 'output/product varieties' for clarity.

[10] The outdegree of i gives the number of varieties to which i supplies, i.e., the number of varieties $j \in \{1, .., N\}$ that use variety i as an input. In contrast, the indegree of i is the number of inputs that i itself uses.

This expression can be decomposed into two parts. The first term in (1) gives the contribution of random adoptions of variety i as an essential input. Recall that each newly introduced variety selects m_K essential inputs uniformly at random from the set of all existing t varieties. Hence m_K/t gives the probability that variety i is selected as a possible essential input. Whether or not the new product t ends up sourcing variety i is determined by an adoption decision that we model below in Section 3.2. For now, we take the adoption probability p_K as given and symmetric across all m_K essential inputs.

The second term in (1) relates to the networked adoption of inputs. It gives the probability that variety i is adopted by the new variety t indirectly, i.e., via the linkages of t's essential inputs. To interpret this term, notice that $A \equiv m_K d_i^{out}(t)/t$ is the expected number of randomly drawn essential inputs that in turn use variety i as an input; in other words, A is the expected number of indirect links that lead from product variety t via its essential inputs k to input variety i.[11] Next, $B \equiv m_N/[m_K(p_K m_K + p_N m_N)]$ is the probability of any given variety in t's network neighborhood to actually be 'drawn' by t, i.e., to be examined more closely as a potential input. To see this, note that the new variety t initially draws m_K essential inputs. In turn, in expectation each of these sources inputs from $p_N m_N + p_K m_K$ varieties.[12] Thus, $m_K(p_K m_K + p_N m_N)$ gives the total number of input links of t's essential input suppliers. In other words, it is the size of the network neighborhood that t searches for potential input varieties. Since t draws m_N (potential) inputs from this network, B is the probability of an input from the network to be drawn. Note that the same input i can show up several times in t's network neighborhood – via different essential inputs. In our car examples, both body and wheels (essential inputs) may use aluminum (network input). This is reflected in the multiplication $A \cdot B$ – the (expected) number of links in t's network neighborhood leading to i, times the probability of any such link to be considered by t as a potential input. Finally, p_N is the probability that an input that has been selected by t as a potential input will actually be adopted.[13] Altogether, the second term in equation (1) thus captures the odds of

[11] To see this, note that the probability that a randomly drawn essential variety k itself sources inputs from variety i is $d_i^{out}(t)/t$, i.e., the number of varieties that i supplies to, divided by the overall number of varieties in the economy in period t. In addition, m_K is the number of such random draws of essential inputs.

[12] This expression also corresponds to the expected indegree, which is the same across varieties in our setup. As for p_K, we take p_N as given for now and model the adoption decision in Section 3.2.

[13] A simple numerical example can provide further illustration: suppose that producer t draws $m_K = 5$ essential inputs, and that the average indegree is 10. Then the size of t's network neighborhood is 50, i.e., there are 50 links leading to further input varieties via t's essential input suppliers. Assume that t decides to closely examine 10 of these input varieties. Then the chance of any input variety from the network neighborhood to be drawn is $B = 0.2$. Next, suppose that input i is extremely prominent, being used by 10% of all varieties. Then $d_i^{out}(t)/t = 0.1$, and $A = 5 \cdot 0.1$ is the expected number of indirect links from t to i, given that t draws 5 essential inputs. Consequently, the chance of i to be drawn by t for closer examination is $A \cdot B = 0.1$. Finally, if t actually adopts half of these potential network inputs, then i has a 5% chance of being adopted by t as a result of the latter's network search.

i being adopted by the new variety t via indirect linkage routes. Importantly, if i already features as an input of a large number of varieties (high $d_i^{out}(t)$), then it is more likely that the new variety also adopts it. This is the core of our mechanism.

Given our setup above, we can characterize the distribution of outdegrees at any time t by means of a mean-field approximation of (1), as in Jackson and Rogers (2007). The mean field approximation is derived by taking a continuous time version of the law of motion in equation (1) where all actions happen deterministically at a rate proportional to the expected change. To do this, let $r \equiv \frac{p_K m_K}{p_N m_N}$ be the ratio of essential inputs to the number of network inputs. In addition, denote by $m = p_N m_N + p_K m_K$ the expected number of inputs adopted by variety t. Then, the following proposition is immediate from Theorem 1 in Jackson and Rogers (2007):

Proposition 1. *In the mean-field approximation of equation* (1)*, the variety outdegree distribution has a cumulative distribution function given by* $F_t(d^{out}) = 1 - \left(\frac{rm}{d^{out}+rm}\right)^{1+r}$ *at any time* t*.*

The proof follows immediately from Jackson and Rogers (2007) and is omitted here.[14] For large d^{out} relative to rm, this approximates a scale free distribution with a tail parameter given by $1 + r = \frac{m}{p_N m_N}$. That is, as the number of network inputs grows large relative to the number of essential inputs, the outdegree distribution of varieties approaches a power law.

3.2 Input Adoption Decision

In the following, we describe the input adoption decision in detail. A new variety producer t decides which inputs to adopt from the set of essential inputs, K_t, and from the set N_t of potentially useful inputs that were identified during the network search stage. The adoption decision is driven by a trade-off between two forces. On the one hand, a producer benefits from a larger set of input varieties, as in standard endogenous growth models in the spirit of Romer (1990). On the other hand, there is a variety-specific customization cost for each adopted input. To model the input adoption decision, we introduce a production function that uses other varieties as intermediates together with labor. Thus, the underlying production structure is a network of linkages across varieties. We focus on a partial equilibrium analysis and illustrate the tradeoff that governs the adoption decision in the symmetric case.

Variety Production

We begin by clarifying notation. We use k to denote elements of the set of essential inputs K_t, and n for network inputs in N_t. Note that both these sets represent *potentially* used inputs. Let

[14] The quality of this mean field approximation can be checked against simulations of the original law of motion. As Jackson and Rogers (2007) show, the mean field result above accords well with simulated distributions of the actual process.

$\widehat{K}_t \subseteq K_t$ and $\widehat{N}_t \subseteq N_t$ be the subsets of essential and network inputs, respectively, that are actually adopted. In the following, we model the decision of a new variety producer t who decides which inputs to adopt.

Each product variety t uses other varieties as intermediate inputs. Their quantities are denoted by x_{tk} and x_{tn} for essential and network inputs, respectively. For illustration, we keep the sets of essential and network inputs separate in the production function, by assuming that they enter two different composites. Inputs of each category enter production as substitutes with elasticity $\epsilon > 1$, so that the corresponding composites are given by:

$$\mathbf{X}_t^K = \left(\sum_{k \in \widehat{K}_t} x_{tk}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \quad \text{and} \quad \mathbf{X}_t^N = \left(\sum_{n \in \widehat{N}_t} x_{tn}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \tag{2}$$

In order to adopt an input, it must be customized at a cost that is specific to each product-input pair. For example, customizing a light sensor for a car is different from customizing a light sensor for an outdoor lamp, and both in turn are different from customizing a rear view camera for a car. We denote this product-input specific customization cost by $c_{t,k}$ and $c_{t,n}$ for essential and network inputs, respectively. To simplify the analysis, we assume that the customization cost is negligible for essential inputs, so that $c_{t,k} = 0, \ \forall k \in K_t$. This reflects our interpretation that a variety's essential inputs are fundamental parts whose integration is standardized, such as wheels or an engine for a car.[15] Because the input composites in (2) feature returns to the number of varieties, the optimal decision for the producer of t is to adopt *all* essential inputs $k \in K_t$.[16]

On the other hand, adopting network inputs is subject to the customization cost $c_{t,n} > 0, \ \forall n \in N_t$. These are calculated as $c_{t,n} = b \cdot r_{t,n}$, where $b > 0$ and $r_{t,n}$ is uniformly distributed over the unit interval. The total cost of adopting a subset $\widehat{N}_t$ of these inputs is given by

$$C_t = \sum_{n \in \widehat{N}_t} c_{t,n} \tag{3}$$

We assume that the customization cost is paid in units of t's output, y_t, in every period of production.[17] This ensures that our results are not driven by scale effects.[18] We can now specify

[15] We build on this notion below when aggregating varieties into sectors.

[16] To see this, note that in the symmetric case, $\mathbf{X}_t^K = \widehat{K}_t^{\frac{1}{\epsilon-1}} \cdot \left(\widehat{K}_t \bar{x}_{Kt} \right)$, where $\bar{x}_{Kt}$ is the quantity used of each essential input. Thus, the more essential inputs are adopted (higher $\widehat{K}_t$), the larger is $\mathbf{X}_t^K$, for any given total amount of essential inputs used ($\widehat{K}_t \bar{x}_{Kt}$).

[17] Thus, C_t can be thought of as annualized customization cost, paid in units of output.

[18] In contrast, if C was a fixed cost, higher demand for a given variety would also lead it to adopt more inputs. This would render the basic structure of our model untractable. In addition to ensuring tractability, this setup is also in line

the variety production function. The two input composites $\mathbf{X}_t^K$ and $\mathbf{X}_t^N$ enter in a Cobb-Douglas fashion, in combination with labor, l_t.[19] For a given (annualized) input customization cost C_t, the output of variety t is given by:

$$y_t = \frac{A_t}{1+C_t} \left(\mathbf{X}_t^K\right)^{\alpha} \left(\mathbf{X}_t^N\right)^{\beta} l_t^{1-\alpha-\beta} \tag{4}$$

where A_t is the productivity draw of producer t. Note that $C_t < 1$ must hold, and that C_t can be interpreted similar to a tax on output, used to cover the initial adoption cost.[20]

Optimization and Input Adoption

A variety producer t solves the cost minimization problem associated with (4), by choosing the set of network inputs $\widehat{N}_t$, as well as the quantity of each input. We begin by analyzing the latter. For given sets K_t and $\widehat{N}_t$, the optimal choice of input quantity x_{ik} and x_{in} in the two aggregates in (2) yields the corresponding price indexes[21]

$$\Phi_t^K = \left(\sum_{k\in K_t} \phi_k^{1-\epsilon}\right)^{\frac{1}{1-\epsilon}} \quad \text{and} \quad \Phi_t^N = \left(\sum_{n\in \widehat{N}_t} \phi_n^{1-\epsilon}\right)^{\frac{1}{1-\epsilon}} \tag{5}$$

where ϕ_k and ϕ_n are the prices of essential and network inputs, respectively. Labor l_t is also chosen optimally, taking the wage w as given. The marginal cost of producing variety t is then

$$MC_t = \frac{1+C_t}{A_t} \left(\frac{\Phi_t^K}{\alpha}\right)^{\alpha} \left(\frac{\Phi_t^N}{\beta}\right)^{\beta} \left(\frac{w}{1-\alpha-\beta}\right)^{1-\alpha-\beta} \tag{6}$$

This expression holds for a *given* set of adopted network inputs $\widehat{N}$. Next, we obtain the optimal set of network inputs, by collecting the terms in (6) that depend on this choice, C_t and Φ_t^N, and

with our technological interpretation that once a variety has chosen its inputs, these are stable over time – that is, a variety is defined by its input use.

[19] Thus, the two input composites are gross complements. This assumption does not affect our qualitative results – we could alternatively assume that the two composites are substitutes, or we could include all inputs in one aggregator. The advantage of our formulation is that we can separate essential inputs and network inputs in a straightforward fashion.

[20] The optimization problem described below ensures this condition as long as at least one network input n has an associated customization cost $c_{t,n} < 1$.

[21] We use the notation K_t rather than $\widehat{K}_t$ to underline that all essential inputs are adopted.

substituting from (3) and (5):

$$\widehat{N}_t^* = \underset{\widehat{N}_t \subseteq N_t}{\arg\min} \left\{ \left(1 + \sum_{n \in \widehat{N}_t} c_{t,n} \right) \left(\sum_{n \in \widehat{N}_t} \phi_n^{1-\epsilon} \right)^{\frac{\beta}{1-\epsilon}} \right\} \tag{7}$$

If the set N_t has many elements, this is a complex combinatorial problem that must be solved numerically. Note that for each potential input variety n in t's network neighborhood, a lower price ϕ_n makes adoption more likely. Thus, technological progress in variety production can raise the rate of adoption, by lowering the input price. We will test this prediction in our empirical analysis. In the following, we illustrate the adoption decision by focusing on the simplified symmetric case.

Symmetry and Illustration of the Adoption Decision

To simplify the analysis, we use the fact that our model implies – in expectation – symmetry across varieties. First, let each variety have the same technology draw $A_t = A$ and assume that demand is such that the price of each variety is a constant markup over its marginal cost.[22] In addition, in expectation each variety uses the same number of essential inputs, m_K, and it draws the same number of potentially useful network inputs, m_N. What remains to be shown for the symmetric equilibrium is that each variety also *adopts* – in expectation – the same number of network inputs.

Adoption costs are also symmetric in expectations, but their realizations vary across the input varieties in N_t. We can thus rank the m_N network inputs in N_t by their adoption costs, such that $c_{t,1} < c_{t,2} < ... < c_{t,m_N}$. Because customization costs are uniformly distributed, the ordered draws $n = 1, ..., m_N$ will lie (in expectation) on the line $c_{t,n} = b \cdot \frac{n}{m_N}$. Let $\widehat{m}_N \leq m_N$ denote the number of adopted inputs (i.e., the size of the set $\widehat{N}_t$). Then the total cost of customization is given by $\sum_{n=1}^{\widehat{m}_N} c_{t,n} = \frac{b}{m_N} \frac{\widehat{m}_N(\widehat{m}_N+1)}{2}$, which is increasing and convex in $\widehat{m}_N$. In expectation, this customization cost function is the same for each new variety t. Consequently, each new variety is expected to adopt the same number of inputs from its network environment. In other words, the indegree is the same for all varieties. Thus, in expectation our model features a symmetric equilibrium with all new varieties facing the same marginal cost in (6) and therefore charging the same price. Note, however, that variety producers use different *sets* of inputs. Thus, the outdegree may be asymmetric – some varieties are more popular suppliers than others. Nevertheless, the total demand for an input affects neither its pricing nor its own adoption of inputs. Consequently, in our setup, symmetry of prices is compatible with asymmetry in the number of forward linkages.

Under symmetry of prices ($\phi_n = \phi, \ \forall n$), and given the above ranking of customization costs,

[22] This follows if we assume that all varieties are aggregated into a final good with elasticity of substitution ϵ. Then both final and intermediate demand for all varieties imply the profit-maximizing markup $\epsilon/(\epsilon - 1)$.

(7) simplifies to:

$$\widehat{m}_N^* = \underset{\widehat{m}_N \leq m_N}{\arg\min} \left\{ \left(\frac{1}{\widehat{m}_N} \right)^{\frac{\beta}{\epsilon-1}} + \frac{b}{2m_N} \frac{\widehat{m}_N(\widehat{m}_N+1)}{(\widehat{m}_N)^{\frac{\beta}{\epsilon-1}}} \right\} \phi^{\beta} \tag{8}$$

The first expression in (8) is decreasing in $\widehat{m}_N$, while the second expression is increasing if $\beta < 2(\epsilon-1)$.[23] This delivers the U-shape shown in Figure 5. To illustrate the intuition for this functional form, the ranking of network inputs by their (randomly drawn) customization costs is crucial. When few inputs are adopted (low $\widehat{m}_N$), customization costs of these low-ranked inputs are small, and therefore the input variety effect à la Romer (1990) dominates. For higher $\widehat{m}_N$, customization costs for each additional adopted input are larger, outweighing the input variety effect. Thus, production costs become increasing in $\widehat{m}_N$. The optimal number of adopted network inputs, $\widehat{m}_N^*$, corresponds to the minimum of the U-shaped curve given by (8).

[Insert Figure 5 here]

Note that our analysis in the symmetric case endogenizes the probability p_N of adopting network inputs, which we took as given in (1). Each new variety draws m_N network inputs, and according to (8), it will adopt $\widehat{m}_N^*$ of these. The likelihood of adoption is thus a-priori the same for any network input in the set N_t, and it is given by $p_N = \widehat{m}_N^*/m_N$. Finally, because of price symmetry, a variety producer j uses the same amount of each input variety i, conditional on this input being used ($b_{ij} = 1$). Thus, the corresponding value of the input purchase, v_{ij}, is proportional to the binary variable b_{ij}. This becomes important below when we aggregate our model to the sector level: our variety level predictions are derived for the *unweighted* directed network (based on binary b_{ij}), while input-output data deliver a *weighted* (value-based) network. Due to the proportionality, variety-level predictions hold at the sector level.

3.3 Sector Level Implications

The model of networked input adoption laid out above is defined at the variety level. Yet, data on variety level input-output networks are not available. The closest data counterparts with a wide coverage are for sector-level input-output networks. Thus, in order to render the underlying model of network formation testable, we now explore its sectoral implications.

[23] For example, suppose that the overall expenditure share for intermediate inputs is 0.5, and that half of these are network inputs. Then $\epsilon > 1.125$ will ensure that the second expression in (8) is decreasing in $\widehat{m}_N$. As a comparison, the average elasticity of substitution reported by Broda and Weinstein (2006) is 4.

Aggregation of Varieties into Sectors

We start by defining how varieties are assigned to sectors in the context of our model. We employ a principle of similarity of inputs. As a result, sectors are composed of varieties that share similar production processes, i.e., varieties that rely on similar input bundles. This input-based approach is also a guiding principle of actual sectoral classification systems like NAICS.[24] To capture this notion, we define a binary baseline vector μ_{s_j} that defines a sector s_j based on its inputs. This can be thought of as a blueprint for the typical inputs used by varieties in sector s_j. For example, the car sector may be represented by a baseline vector μ_{s_j} with unit entries in 'glass windows', 'engine', and 'wheels'. The vector μ_{s_j} can be thought of as the classification scheme for new varieties. Each variety is then classified into the sector whose μ_{s_j} has the maximum overlap with the variety's list of *essential* inputs.[25] In other words, a variety's essential inputs are compared to the typical inputs used by all sectors in the economy, and it is then classified into the most similar one. The following definition formalizes this principle:

Definition 1. (Definition of a Sector): *At time t, a sectoral classification system is a partition of the set of existent varieties into J sectors. Each sector s_j, with $j = 1, ..., J$, is defined by a t-dimensional binary vector, μ_{s_j}, with a total of x ones and $t - x$ zeros, with unit entries in the vector being elected at random. Each existent variety is assigned to a sector by finding the sector s_j that maximizes the overlap between that variety's binary vector of essential inputs and the vector μ_{s_j}. Any new variety introduced at time $t + 1$ is classified into a sector in the same way.*

Note also that we are allowing for overlap among sectors, in that different sectors can share some elements across their baseline vectors. For example, 'tires' can be represented in both the bicycle and car sectors. Note that this definition induces a sectoral input-output network of dimension $J \times J$, where nodes are now sectors and directed edges, $a_{s_i s_j}$, represent intersectoral input flows from sector s_i to sector s_j. According to our definition, these directed edges reflect varieties which have been classified into sector s_j and source inputs from varieties classified into sector s_i.[26]

[24]For example, the Bureau of Labor Statistics provides a detailed explanation of this production-based principle: "Industries are classified on the basis of their production or supply function – establishments using similar raw material inputs, capital equipment, and labor are classified in the same industry" (Murphy, 1998, p.44). This refers to the more recent NAICS system. The previous SIC classification system used information on inputs employed but also took into account the uses/demand for the good produced.

[25]This notion of overlap can be made formal by use of the Hamming distance between two binary vectors of the same length. This distance gives the number of elements by which the two binary vectors differ. Thus, we classify a given variety into the sector s_j whose baseline vector μ_{s_j} has the minimum Hamming distance to this variety's essential inputs.

[26]For a fixed number of sectors J, as t becomes large, eventually all sector pairs will exhibit non-zero flows $a_{s_i s_j}$. We study sector-level adoption, meaning that $a_{s_i s_j}$ goes from zero to positive. We thus implicitly assume that the time t input-output network is sparse, i.e., that many $a_{s_i s_j}$'s are zero. In addition, note that economies with zero $a_{s_i s_j}$

Sector-Level Predictions

We now turn to the evolution of the sector-level input-output network over time. At the variety level, the key mechanism of network formation relied on a notion of network proximity: a new variety is more likely to adopt inputs in its network neighborhood, as defined by the set of varieties that supply inputs to the new variety's essential inputs. We now show that such a mechanism is still present under aggregation at the sectoral level. To see this, we first define a sector-level measure of network proximity for any ordered pair of sectors for which there is no input supply relation at time t. This definition exploits variety-level input flows from sector s_i to sector s_j

Definition 2. (Sector-level Network Proximity): *Take any ordered pair of sectors* (s_j,s_i) *such that* $a_{s_i s_j} = 0$ *at time* t. *The network proximity of* (s_j, s_i) *is defined as* $n_{(s_j,s_i)} \equiv \mu'_{s_j}\nu_{s_i}$ *where* ν_{s_i} *is a* $t \times 1$ *vector, where each entry* $\nu_{s_i}(v)$ *gives the number of varieties from sector* s_i *that are sourced as inputs by variety* v, *for* $v = 1, .., t$. *We say that sector* s_j *is closer to* s_i *than* $s_{j'}$ *if* $n_{(s_j,s_i)} > n_{(s_{j'},s_i)}$.

This definition states that sector s_i is closer to s_j if varieties from s_i are used more frequently as inputs by varieties that define sector s_j. That is, $n_{(s_j,s_i)}$ gives the number of varieties in sector s_i that are sourced as inputs by varieties which appear in the baseline vector of sector s_j.[27] Next, we use this definition to aggregate varieties to the sector level. A new variety t will be classified into the sector s_j whose baseline vector is most similar to t's essential inputs. The sector-level network proximity $n_{(s_j,s_i)}$ then tells us how closely we should expect t to be connected to inputs from each sector s_i. Intuitively, if t is classified into j, it must have a relatively large number of essential inputs that are also present in s_j's baseline vector μ_{s_j}. Thus, t must also have many input links in common with the varieties in μ_{s_j}. This is the proximity dimension that $n_{(s_j,s_i)} \equiv \mu'_{s_j}\nu_{s_i}$ exploits. Given this definition, the following proposition shows that the network proximity mechanism underlying the variety level model is still present when we aggregate varieties into sectors.

Proposition 2. *Take any two sectors* s_j *and* $s_{j'}$ *that previously did not source inputs from sector* s_i, *i.e.,* $a_{ij} = a_{ij'} = 0$ *at* $t-1$. *Then if at time* $t-1$ *sector* s_j *is closer to* s_i *(i.e.,* $n_{(s_j,s_i)} > n_{(s_{j'},s_i)}$*),* s_j *will be more likely to adopt an input from* s_i *at* t.

We provide a formal proof in the appendix. Here, we briefly describe the intuition. First note that any new input linkages at period t must be due to the new variety t; all pre-existing varieties do not change their linkage structure. Whether t links s_j and s_i depends on (i) whether t is classified

can be maintained even for large t if the sectoral classification system is expanded by raising J, i.e., by refining the sectoral detail.

[27] Note that this proximity definition need not be symmetric, i.e., generically $n_{(s_j,s_i)} \neq n_{(s_i,s_j)}$, as is standard for network distance metrics in the context of directed graphs.

as an element of sector s_j, and (ii) whether it then sources input(s) from sector s_i. The proof links both steps by following the classification scheme for sectors described above: The new variety t randomly draws a set of essential inputs. It is then classified into the sector s_j that has the closest overlap with these essential inputs. Thus, the fact that t is sorted into sector s_j tells us that it shares more essential inputs with varieties in s_j than with varieties in any other sector $s_{j'}$. This is criterion (i). Criterion (ii) then incorporates new link formation via the network neighborhood of t's essential inputs. If many of these link to sector s_i, t is more likely to source from s_i.[28] Finally, combining (i) and (ii), if t is classified into a sector s_j that has many indirect input linkages to s_i, t is expected to itself have such indirect linkages to s_i; and these in turn raise the probability that t directly adopts inputs from s_i. Summing up, since a-priory t is equally likely to 'fall' into any sector, the sector s_j closest to s_i (among those that are not yet directly linked to s_i) is most likely to establish a new link to s_i.

Having established that the key network proximity mechanism holds at the sectoral level, we now characterize the size distribution of links. In particular, we are interested in understanding whether our variety level model, when aggregated to the sectoral level, can generate the fat tailed behavior of sectoral outdegrees emphasized in Acemoglu et al. (2012).

To do this, first note that the induced sectoral level network consists of weighted links across sectors, reflecting the number of existing varieties at time t that are both: (i) classified in the same sector s_j and (ii) source as inputs varieties from a given sector s_i. Thus, sector-level input flows $a_{s_i s_j}$ are given by $a_{s_i s_j} \equiv \sum_{i \in s_i} \sum_{j \in s_j} v_{ij}$, where v_{ij} denotes the sales of input variety i to product variety j. In turn, this implies that sector s_i's total sales, i.e., its (weighted) outdegree, are $d^{out}_{s_i} \equiv \sum_{j=1}^{J} a_{s_i s_j}$. Having established this notation we can move on to the following proposition:

Proposition 3. *If the variety-level outdegree distribution at time t is power law distributed, so is the distribution of sectoral outdegrees.*

In the following, we provide a sketch of the proof; for a formal proof see the Appendix. The proof of Proposition 3 relies on the fact that the sum of a finite number of power law distributed random variables is itself a power law random variable. It follows two steps. First, we show that for a large number of varieties t relative to the number of sectors s, each sector at time t contains the same number of varieties in expectations. We then sum across the number of varieties in each sector to prove that the sectoral outdegree distribution is power law distributed.

[28] More generally, the new variety t can form links to inputs in sector s_i directly – drawing $i \in s_i$ as an essential input – or indirectly, via its network of essential inputs. Regarding the former, this initial draw is symmetric across all existing inputs. Thus, it does not differentially affect link-formation across sectors. The proof therefore focuses on the adoption via the network of essential inputs.

First, from the proof of Proposition 2 recall that, ex-ante, the probability of any new variety t being classified into a given sector is the same across sectors. This follows immediately from the joint assumption that both the ideal varieties defining sectors and the set of essential inputs are drawn uniformly at random from the set of existing varieties. Now note that, if this is the case, the expected number of varieties classified into any given sector at time t is also the same across sectors, and it is given by $\frac{t}{J}$, where J is the total number of sectors. If t is much larger than J, the law of large numbers implies that the actual number of varieties classified in each sector at time t is the same across sectors.

Second, under the assumption of price symmetry, the sectoral (weighted) outdegree is proportional to the number of varieties to which a given sector s_i supplies inputs at time t, where the constant of proportionality is given by the price ϕ.[29] Given the observation that for $t >> J$, the number of varieties in any sector is given by $\frac{t}{J}$, a sector's outdegree is simply given by the sum of the variety level outdegree across $\frac{t}{J}$ varieties. Thus, under the assumption that at time t the variety level outdegree distribution is power law distributed, a sector's weighted outdegree is given by the (finite) sum of $\frac{t}{J}$ power law distributed variables. Since power law variables are stable upon aggregation, it is then immediate that a sector's outdegree is itself power law distributed with the same tail exponent as the variety-level outdegree distribution (for a formal proof see Jessen and Mikosch, 2006, Lemma 3.1).

4 Empirical Framework and Data

In this section, we take the model's predictions to the data. While the core mechanism works at the variety level, our aggregation results in Proposition 3 allow us to employ sector level data. We use US input-output benchmark tables between 1967 and 2002 (at the 4 digit level) and track input adoption over time. We then ask whether initial network proximity – measured by existing input linkages – predicts subsequent input adoption. We proceed as follows: we first introduce our measure of network proximity. Second, we describe our data and discuss the definition of adoption in the context of input-output tables. Finally, we present empirical results analyzing both the time to adoption after 1967 and the likelihood of adoption in any given benchmark year. Throughout this section, we use j to denote the input-using (adopting) sector, and i for the input-producing sector.

[29] The weighted outdegree refers to values of input flows, while our variety-level prediction are based on binary (unweighted) input links. Because of price symmetry, product variety j spends the same amount for each input variety i that it uses (see the discussion at the end of Section 3.2). Thus, the overall value of input varieties sold (outdegree) or used (indegree) by a sector is proportional to the underlying *number* of input varieties.

4.1 Network Proximity

When aggregated to the sectoral level, our model predicts that sector j is the more likely to adopt input i the more closely j is already related to i via indirect network connections. In the following, we use a standard measure of network distance that captures this notion. It builds on the hypothesis that sectors trading inputs more intensively are 'closer' in the technology landscape.[30] Crucially, the distance measure can also be calculated if there is no direct path linking two sectors – in this case we compute the shortest path via intermediate steps.

Formally, we define a direct-requirements input-output matrix Γ where each element Γ_{ij} represents the cost share of input i in the total intermediate input expenditures of sector j. If Γ_{ij} is non-zero, we define the distance from j to i as $d_{ij} = \frac{1}{\Gamma_{ij}}$. Thus, the more important input i is in the production of j, the closer is d_{ij} to 1 (the minimum possible distance between two sectors). The case $\Gamma_{ij} > 0$ holds if a direct connection between i and j exists, i.e., if j has already adopted i. However, since we study adoption, the relevant starting point is $\Gamma_{ij} = 0$.

Provided that j *indirectly* sources inputs from i – via its network of suppliers – we define the distance d_{ij} as the sum of the distances along the shortest path that connects i and j. For example, if j uses input k, which in turn sources inputs from i, then $d_{ij} = d_{ik} + d_{kj}$.[31] If there exist more than one such paths linking j and i, then d_{ij} is the minimum distance path, i.e. the directed path between the two nodes such that the sum of the weights of its constituent edges is minimized. This shortest path algorithm yields distances between any two sectors in the economy.

4.2 Data and Main Variables

In the following, we describe our dataset and the derivation of our main variables. We use y (years) to denote the time dimension, in order to avoid confusion with the variety index t above.

Input-Output Data

We calculate the measure of network distance d_{ij}, using the Bureau of Economic Analysis (BEA) Benchmark Input-Output Use Tables. The BEA provides U.S. input-output (I-O) data at the 4-digit SIC level in 5-year periods (benchmark years) between 1967 and 2002. Following Carvalho (2010)

[30] Note that our model makes two simplifying assumptions. First, local search occurs only at the level of two degrees of separation (i.e., across direct neighbors of 'parents'). Second, the model emphasizes the *number* of (indirect) routes, abstracting from the intensity of linkages. In the data, however, adoptions can occur between sectors that are initially more than two nodes apart. Also, the intensity of linkages (input shares) is not symmetric in the data. Our empirical measure of distance captures both these features.

[31] See, for example, Ahuja, Magnanti, and Orlin (1993) or Jackson (2008) for a review of distance and shortest path measures in networks. Note that, in principle, the distance measure can also be calculated in the opposite direction – looking for the shortest path from inputs i to sector j, d_{ji}. Our use of d_{ij} reflects the notion of distance implied by our model, where the new variety t purchases essential inputs k. The extent to which the latter are connected to i determines the likelihood of t adopting i.

and Acemoglu et al. (2012), we view the input-output matrix as a network of input-flows, where each sector is a node, and each input-supply relationship is a (weighted) directed edge linking two nodes.

For some sectors, the level of aggregation or coverage changes over time. We account for this by aggregating sectors, and match the resulting I-O panel to the Annual Survey of Manufacturing (ASM) 1987 SIC classification.[32] In 1997, the BEA changed the I-O classification from SIC to NAICS. While the Census Bureau provides a correspondence, the match is imperfect for many sectors at the 4-digit level. To make sectors comparable beyond the last SIC-based I-O table in 1992, we employ the following procedure: (i) if several NAICS sectors match a single SIC sector, the former are aggregated; (ii) if several SIC sectors were merged into one NAICS sector in 1997, industry-commodity specific shares from the 1992 I-O table are used to disaggregate NAICS into the corresponding SIC components.[33] The switch to NAICS also reclassified products into new sectors, and the correspondence assigns these in part to existing SIC sectors. This creates events that look like adoption in 1997.[34] To avoid that this affects our results, we exclude new linkages formed in 1997 in our baseline analysis. Nevertheless, our robustness analysis shows that most results go through even if we add the noisy 1997 data.

Overall, our approach to making sectors comparable yields a coherent set of 358 sectors for all I-O benchmark years between 1967 and 2002. For each sector-input pair, we calculate our central explanatory variable: network distance in 1967, d_{ij}^{67}. To identify the minimum distance path between i-j pairs, we use a standard Dijkstra's shortest path algorithm (see for example Ahuja et al., 1993).

Input Adoption and Time to Adopt

We define input adoption as an event in a given year y, where a sector j begins to use an input i. We say that j has adopted i in y if it has not used the input prior to year y, and begins to purchase a positive amount of the input in y. Formally, the indicator variable for adoption in year y is thus defined as:

$$A_{ij}(y) = \begin{cases} 1, & \text{if } \Gamma_{ij}(y) > 0 \text{ and } \Gamma_{ij}(y') = 0, \ \forall y' < y \\ 0, & \text{otherwise} \end{cases} \tag{9}$$

[32] For a detailed description of this methodology see Voigtländer (2014). One example are paper mills (SIC 2621) and paperboard mills (SIC 2631). Both are reported separately in the I-O data before 1987, but aggregated to one sector thereafter. We treat these data as one sector, 'paper and paperboard mills' over the full sample period.

[33] The original NAICS-SIC correspondence is available at http://www.census.gov/epcd/www/naicstab.htm. The extended correspondence that includes industry-commodity specific weights is available upon request from the authors.

[34] The 2002 I-O data, on the other hand, are directly comparable with their 1997 counterpart, so that we can clearly identify adoption events in this year.

Note that this definition yields $A_{ij}(y) = 0$ in the cases of pre-existing links and when an input connection between i and j existed in the past but disappears in y (broken links).

We compute two definitions of adoption, a broad (A_{ij}^{br}) and a narrow one (A_{ij}^{nar}), using 5-year intervals corresponding to IO benchmark years. A_{ij}^{br} requires that i has not been used in $y-5$, and is used in y. Therefore, the broad definition potentially also captures cases where inputs are adopted and then dropped again.[35] Many of these short-term adoption events are probably noise, but some may also reflect actual attempts to integrate new inputs. The narrow definition excludes such events, requiring that i be used for at least 10 years after adoption, i.e., in $y+5$ and $y+10$. This comes at the cost of 'losing' adoptions during the last two benchmark years in our sample. We use the broad definition as our main measure and document the robustness of our results using the narrow measure.

Next, we define the time that it takes a given sector to adopt an input:

$$T_{ij} = y_{Adopt} - 1967 \,, \tag{10}$$

where y_{Adopt} is the year in which sector j adopted input i; formally, $A_{ij}(y_{Adopt}) = 1$. Note that this measure is only defined if i) there was no input link between i and j in 1967 ($\Gamma_{ij}(1967) = 0$), and ii) adoption occurred before the end of our sample in 2002. Altogether, there are 128,164 i-j pairs in our dataset. Out of these, 16,684 have $\Gamma_{ij} > 0$ in 1967, which leaves 111,480 possible adoption events. During the subsequent four decades until 2002, we observe 21,161 adoptions in our broad measure and 8,783 in the narrow one.[36]

Sectoral Characteristics

We use sector-level data from the NBER-CES Manufacturing Industry Database, which provides total factor productivity (TFP), output price deflators, wages, value of shipments, and capital stock at the 4-digit SIC level over the period 1958-2005. These data are collected from various years of the Annual Survey of Manufactures (ASM).[37] We use these data to derive control variables for input producing and adopting sectors. Most importantly, we calculate changes in TFP for input producing sectors, $\triangle TFP_i$. We use this to test the prediction that sectors with rapid productivity growth are more likely to be adopted. Since this variable may be endogenous to adoption, we also compute the changes in TFP before 1967, starting from the earliest year for which data is available, 1958. This variable, $\triangle TFP^{58-67}$, strongly predicts TFP growth after 1967.

[35] However, multiple adoption events are excluded because our definition yields $A_{ij}(y) = 0$ if a connection between i and j had existed before.

[36] As discussed above, this excludes 1997 to avoid that adoption events reflect the change from SIC to NAICS in that year.

[37] See Bartelsman and Grey (1996) for a documentation.

5 Empirical Results

In this section, we test our model's main prediction that closer network proximity raises the likelihood of subsequent input adoption. We approach this question in two ways. First, we use a panel approach to show that the probability of adoption of input i by sector j in year y depends on technological distance d_{ij} at $y-5$ (i.e., in the previous I-O benchmark year). Second, we show that in the cross-section of sectoral i-j pairs, adoption tends to happen earlier for smaller initial network distance d_{ij}^{67}. We also show that, in line with our model, more rapid technological progress in an input producing sector raises the odds of adoption.

5.1 Panel Estimation: Probability of Adoption

Does closer network proximity raise the likelihood of input adoption? In the following, we examine this question in the context of a panel in 5-year intervals between 1967 and 2002. For each I-O benchmark year y, we compute our distance measures $d_{ij}(y)$ as described in section 4.2. For all i-j pairs that were not directly connected in any year prior to y, we ask whether the probability of adopting in year y depends on our lagged network distance measure $d_{ij}(y-5)$:

$$Prob\left(A_{ij}(y)=1\right) = g\left(\ln d_{ij}(y-5), X_i(y), X_j(y)\right) \ , \tag{11}$$

where X_i (X_j) are additional controls for the input-producing (adopting) sector, such as changes in total factor productivity or fixed effects. We use log distance to avoid that outliers affect our results disproportionately. The dependent variable in each regression is the indicator $A_{ij}(y)$ as defined in (9).[38] We estimate different functional forms $g(\cdot)$. Given the binary nature of the dependent variable, our main specification is the probit model. We also estimate a linear probability model and a hazard model, finding very similar results.

Main Results

We begin by reporting results for our baseline specification – the Probit model – in columns 1 and 2 of Table 1. The coefficient on network distance is highly significant and negative; thus, lower initial network distance makes adoption more likely. In order to interpret the magnitude of coefficients, we also report standardized coefficients in square brackets for our two main explanatory variables: network distance and TFP in input producing sectors. They show how a one standard deviation increase in the respective explanatory variable affects the probability of adoption. With a standardized coefficient of -2.34 percentage points, the effect of network distance is economically

[38] Note that this definition excludes all (directed) i-j pairs with input flows prior to y. This also implies that upon input adoption in y, the corresponding i-j pair is excluded from the sample in all years $y' > y$.

significant.[39] The coefficient remains unchanged in column 2, which controls for TFP growth over the previous five years in both the input-producing (i) and adopting sector (j). The coefficient on $\triangle TFP_i$ is positive and highly significant, but the magnitude is markedly smaller – with a standardized effect of 0.07 percentage points for an average i-j pair. The differences in magnitude suggests that network proximity is the quantitatively more important driver of pair-specific input adoption. Finally, sectors j that see more rapid TFP growth ($\triangle TFP_j$) are *less* likely to adopt new inputs. This is compatible with Helpman and Trajtenberg (1994), who argue that the actual productivity benefits from a change in production methods may materialize later, so that periods of technology adoption are associated with a temporary slowdown in productivity. However, the effect is quantitatively minuscule, with a standardized coefficient of 0.05 percentage points (not reported in the table).[40]

In columns 3 and 4 in Table 1 we show that our results also hold in a simple linear probability model (OLS). According to the estimate in column 3, a one standard deviation (std) increase of $d_{ij}(y-5)$ raises the probability of adoption throughout the following five years by 1.45 percentage points. The coefficient remains unchanged in column 4, which controls for TFP growth over the previous five years. TFP changes in input producing and adopting sectors have the same sign and significance as in the Probit model, and both remain quantitatively small.

[Insert Table 1 here]

In columns 5 and 6 we estimate a proportional hazard model. The hazard ratio for distance (0.595) implies that as $d_{ij}(y-5)$ increases by one unit, the rate of adoption in any given period will be 0.595 as high as before, i.e., it will decrease by 40.5%. Alternatively, a one std increase in $d_{ij}(y-5)$ reduces the adoption rate by 56.4%.[41] The corresponding standardized *relative* hazard coefficient is -4.2 percentage points, implying that over the entire sample period, a one standard deviation increase in network distance is associated with a -4.2 pp. lower probability of adoption. TFP growth in both input producing and adopting sectors have hazard ratios above 1, indicating that TFP growth is associated with faster adoption. While this confirms the Probit and OLS results for input producing sectors, it contradicts them for TFP growth in adopting sectors. However, the magnitude of both effects remains small, with standardized coefficients in the range of 0.1%. In

[39] The marginal effect implied by the Probit coefficient in column 1 is -0.0147, and the standard deviation of network distance is 1.60.

[40] Our model does not predict a sign for the coefficient on TFP growth in the adopting sector, $\triangle TFP_j$. Empirically, the coefficient is not robust and changes signs in the specifications below. For example, in the Hazard regressions, higher $\triangle TFP_j$ is associated with faster adoption.

[41] The coefficient is ln(0.595)=-0.5196, and thus exp(-0.5196*1.6)=0.436 is the hazard ratio for a one std (1.6) increase in network distance.

sum, the hazard model confirms the economically and statistically significant (negative) effect of network distance on the odds of input adoption, as well as the quantitatively small positive effect of TFP growth in input producing sectors.

Additional controls, sector fixed effects, and robustness

In Table 2 we present alternative specifications and include additional controls, using Probit regressions. Columns 1-3 use our broad measure of input adoption; columns 4-6 use the narrow one, which requires new i-j links to persist for 15 years in order to be counted as adoption. In addition to the broad/narrow categories, the measures of network distance also vary in two additional dimensions: first, columns 2 and 4 exclude input links that are formed between 4-digit sectors within the same 2-digit industry. This reduces the number of adoption events by 8%.[42] Thus, most input adoptions occur across 2-digit sectors. Second, columns 3 and 6 use network distance measured at the beginning of the sample period, in 1967. All regressions in Table 2 control for the level of TFP and employment in adopting (j) and input-producing (i) sectors. Controlling for sector size (employment) captures an important potential confounding factor – that larger sectors may be mechanically more connected and more likely to adopt.

[Insert Table 2 here]

We find that neither the additional controls nor the variations in the network distance measure changes our results. Throughout the specifications, network distance has a strong negative effect on adoption probabilities. For lagged distance, this effect is very similar in magnitude to the results in Table 1 – a one std decline in $d_{ij}(y-5)$ raises the odds of adoption in y by 1.9 percentage points. When using distance in 1967 (columns 3 and 6), a one std reduction in d_{ij}^{67} (0.65) raises the probability of adoption by approximately 1.2 percentage points. The somewhat smaller estimate is probably due to the fact that d_{ij}^{67} becomes an increasingly more imprecise measure towards the end of our sample period. In line with our model, inputs that are produced more efficiently (higher TFP_i) are more likely to be adopted. In our baseline specification (col 1), a one std increase in TFP_i raises the adoption probability by 0.6 percentage points.[43] On the other hand, the coefficients on efficiency of the adopting sector have ambiguous signs and are mostly insignificant. Finally,

[42] For the broad (narrow) measure, we count 21,161 (8,783) input adoption events in our sample (excluding 1997), and this number declines to 19,498 (8,111) when excluding adoption events within 2-digit industries.

[43] This confirms our previous finding that network proximity is the dominant effect. The difference in magnitudes is even more striking for the narrow definition of adoption: the results in col 4 imply that a one std decrease in $d_{ij}(y-5)$ (increase in TFP_i) raises the odds of adoption by 1.8 (0.2) percentage points. Short-run *changes* in TFP ($\triangle_5 TFP_i$) do not have a clear additional impact on adoption – the corresponding coefficient signs are ambiguous. And even for the narrow definition of adoption, where the coefficients are positive and significant, the magnitude is small (with a one std increase in $\triangle_5 TFP_i$ leading to a rise in the adoption probability by 0.3 p.p.).

sector size (measured by employment) is associated with both higher probability of adopting and being adopted.

Is the observed effect of network distance on the adoption rate merely driven by unobserved sectoral characteristics? For example, more 'dynamic' sectors may be more central in the input-output network and also adopt new inputs more frequently. In Table 3 we address this issue by including fixed effects for input-producing and input-using sectors, in addition to benchmark year dummies. Both significance and magnitude of the coefficient on network distance are unchanged. The same is true for the coefficient on TFP in input producing sectors (TFP_i).[44] In the presence of fixed effects, the positive coefficient on TFP_i means that adoption of an input is more likely in periods when its TFP is high, relative to its own average *and* relative to the average across all other sectors in the same year. In other words, adoption is more likely in periods when the input-producing sector performs particularly well. As before, TFP in adopting sectors shows no clear relationship with the likelihood of input adoption. Finally, the relationship between input adoption and employment is now ambiguous for input producing sectors (i), and less robust than above for adopting sectors (j). Finally, Panel B of Table 3 documents very similar results when we restrict adoption events to purchases above \$1 mio. This ensures that minor transactions in input-output tables do not affect our results.

[Insert Table 3 here]

5.2 Cross-Sectional Estimation: Time to Adoption

In the following, we analyze how initial network distance in 1967 affects the time that it takes until a sector j adopts an input i, T_{ij}. This is conditional on adoption being observed by the end of our sample period in 2002. We run the following regression:

$$T_{ij} = \beta \cdot d_{ij}^{67} + \gamma \cdot \triangle Efficiency_i + \delta_i + \eta_j + \varepsilon_{ij} \ , \tag{12}$$

where d_{ij}^{67} is network distance in 1967, and $\triangle Efficiency_i$ denotes the (average annual) change in efficiency in the input producing sector.[45] To proxy for efficiency, we use TFP changes and changes in input prices between 1967 and the year of adoption, as well as pre-1967 TFP changes in the input producing sectors. Finally, δ_i and δ_j are input producing and adopting sector fixed effects, respectively.

[44] The implied marginal effects are also very similar to those documented above: a one std decrease in $d_{ij}(y-5)$ (increase in TFP_i) raises the odds of adoption by 1.9 (0.5) percentage points

[45] By using average annual changes, we avoid that later adoption is mechanically associated with higher efficiency gains, because technology in general advances more over longer horizons.

Table 4 reports the results, using OLS regressions. We use fixed effects for input adopting sectors (δ_j) throughout, capturing the large degree of heterogeneity across sectors. As for input producing sector fixed effects, there is a tradeoff. On the one hand, some sectors are more central in the network than others, which we expect to raise their likelihood of being adopted. Using fixed effects δ_i will absorb this variation, which may attenuate our results. On the other hand, there are many other potential sector-specific features that may confound our results; including δ_i controls for those that are time-invariant. In practice, our results are robust to either specification: col 1 does not include δ_i, while all other specifications in Table 4 do so. The coefficient on network distance is actually stronger when including δ_i, which is probably due to the substantially improved fit of the regression (the R^2 increases from 0.19 in col 1 to 0.73 in col 2). Columns 3 and 4 show that our results are also robust to excluding adoptions that occurred in 1972, as well as to including 1997 (when the IO tables shifted from SIC to NAICS). Finally, excluding adoptions that occurred within 2-digit industries (col 5) and using the narrow definition of adoption (col 6) also yields similar estimates.

In terms of magnitude, a one std decrease in d_{ij}^{67} reduces the time to adopt by 2.14 years, while a one std increase in $\triangle TFP_i(1967 - y_{adopt})$ reduces time to adopt by 2.30 years. In contrast to our results for adoption probabilities in 5-year intervals, the cross-sectional results on time to adopt exploit long-term changes in TFP. These turn out to be a quantitatively meaningful predictor. To put the estimates in context, the average time to adopt (conditional on adoption occurring prior to 2002) in our sample is 16.7 years. Thus, both network distance and TFP growth in i have standardized effects of an approximately 15% reduction in adoption time.

[Insert Table 4 here]

Table 5 provides additional results for time to adoption. In col 1 we control for the change in input prices between 1967 and the year of adoption. As expected, the less prices rise (or the faster they decline), the shorter is the time to adoption. The standardized effect is 4.0 years, which is stronger than the one for TFP changes. Next, in col 2 we use TFP growth prior to 1967. Focusing on historical efficiency growth in i addresses the possibility of reverse causality, i.e., that firms in sector i may anticipate the adoption of i and thus innovate, rather the other way around. The coefficient on $\triangle TFP_i(1958 - 67)$ is highly significant but muted, with a standardized effect of -0.27 years. To obtain a coefficient estimate that can be more readily compared with our baseline results, and at the same time address the possibility of reverse causality, we employ a 2-stage least square approach. We use pre-1967 TFP growth to predict TFP growth between 1967 and the year of adoption.[46] The first stage has very strong predictive power, with an F-statistic above 800. The

[46]Note that this approach is only feasible if we do not use input producer fixed effects δ_i.

second stage results are shown in col 3: output from sectors that see faster TFP growth is adopted significantly faster by other sectors, with a standardized coefficient of -3.9.

[Insert Table 5 here]

In columns 4-6 we add input producer fixed effects. Col 4 shows the results for input price changes, and col 5 also adds TFP changes.[47] Both are highly significant and have the expected sign. The input price effect is quantitatively stronger, with a standardized coefficient of 5.40, as compared to -2.70 for TFP. This is in line with our model, where input cost is a first-order determined of adoption. Finally, col 6 shows that the input price effect also holds for our narrow definition of adoption. Throughout all specifications in Table 5, network distance in 1967 is strongly associated with a shorter time to adoption, with a standardized effect of approximately 2 years.

6 Firm-Level Evidence

In this section we analyze the relationship between network distance and input adoption at the firm level. We use data from COMPUSTAT, which includes information supply linkages.

7 Conclusion

Input-output linkages have important effects on macroeconomic outcomes. By connecting otherwise independent sectors, these linkages propagate sectoral shocks; if some sectors are prominently linked in the network, idiosyncratic shocks can create aggregate fluctuations (Acemoglu et al., 2012). Intersectoral linkages can also rationalize large productivity differences across countries by amplifying idiosyncratic sectoral distortions (Ciccone, 2002; Jones, 2013). While typically observed at the sectoral level, input-output linkages reflect the flow of products between individual producers, and thus ultimately the underlying technology at the product level. Studying the evolution of the input-output structure is therefore at the heart of technological progress.

We study the mechanism of input link formation both theoretically and empirically. We model the evolution of links at the product variety level in a two-step process, where potentially useful inputs are first identified, and then some of them are adopted. Each variety is characterized by a set of essential inputs, and these are in turn connected to additional products via input-output linkages. We refer to this as the network neighborhood of a new variety. In a first step, a new variety producer identifies potential inputs from its network neighborhood. In the second step, actual adoption is driven by a tradeoff between returns to a larger input portfolio à la Romer (1990),

[47] Price and TFP changes are strongly negatively correlated, with a coefficient of -1.61 (standard error 0.036) after including fixed effects δ_i and δ_j.

and input-specific customization costs. Modeling both the essential inputs and customization costs as random draws, we can build on models of social networks in the spirit of Jackson and Rogers (2007) to study the evolution of links across varieties. We then show that aggregation of varieties into sectors – based on actual rules applied in the construction of U.S. input-output tables – delivers two important predictions. First, the distribution of the sector-level outdegree (forward linkages) follows a power law. Second, input adoption is more likely across pairs of sectors that are initially closer in the input-output network. While the power law is a well-documented feature of the input-output network (Acemoglu et al., 2012), we provide novel evidence for the second prediction of our model.

We use detailed US I-O tables at the 4-digit level between 1967 and 2002 to construct a measure of network distance. We provide strong evidence that closer initial network proximity raises the likelihood of subsequent input adoption. This effect is economically important, with a one-std increase in network distance lowering adoption probability by more than 50%. One obvious concern is that network proximity merely reflects the fact that input-output linkages are clustered around the diagonal, i.e., that sectors tend to use their own output, or output of sectors with similar classification codes. To alleviate this concern, we show that our results are equally strong when we restrict attention to the formation of links outside of 2-digit sectors – i.e., far away from the diagonal. We also shed light on the role of adoption costs of inputs. As a proxy, we use the TFP growth of input-producing sectors. This reflects the idea that faster growth in a sector is typically associated with product innovation. In line with our model's prediction, we show that more rapid technological progress in an input-producing sector raises the probability of its adoption.

Our theoretical and empirical results have important implications for the emergence of central nodes – or General Purpose Technologies – in the input-output network. New products that are used by centrally positioned sectors are themselves more likely to evolve into central positions. This feature differentiates our contribution from previous models on the rise of GPTs, where new technologies are a-priori designated to become GPTs (c.f. Helpman and Trajtenberg, 1998). In contrast, in our framework, only a fraction of new varieties develop into central positions. We abstract from scale effects in order to focus on the network dimension of our model. However, this extension may provide interesting implications for economic growth: since product demand is a major driver of innovation in endogenous growth models, input adoption will foster growth by raising the effective market size. This feedback mechanism may accelerate the evolution of central technologies, by raising their attractiveness as inputs. The same mechanism may be a driver of structural change, by creating asymmetries in productivity growth across sectors (Ngai and Pissarides, 2007).

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The Network Origins of Aggregate Fluctuations. *Econometrica 80*(5), 1977–2016.

Ahuja, R. K., T. L. Magnanti, and J. B. Orlin (1993). *Network Flows: Theory, Algorithms and Applications*. Prentice Hall.

Atalay, E., A. Hortaçsu, J. Roberts, and C. Syverson (2011). Network Structure of Production. *Proceedings of the National Academy of Sciences 108*(13), 5199–5202.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013). The Diffusion of Microfinance. *Science 341*(6144).

Bartelsman, E. J. and W. Grey (1996). The NBER Manufacturing Productivity Database. *NBER Technical Working Paper 205*.

Bigio, S. and J. La'O (2013). Financial Frictions in Production Networks. Working paper.

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

Carvalho, V. M. (2010). Aggregate Fluctuations and the Network Structure of Intersectoral Trade. Working Paper.

Chaney, T. (2013). The Network Structure of International Trade. mimeo. U. Toulouse.

Ciccone, A. (2002). Input Chains and Industrialization. *Review of Economic Studies 69*, 565–587.

Conley, T. G. and B. Dupor (2003). A Spatial Analysis of Sectoral Complementarity. *Journal of Political Economy 111*(2), 311–352.

Conley, T. G. and C. R. Udry (2010). Learning about a New Technology: Pineapple in Ghana. *American Economic Review 100*(1), 35–69.

di Giovanni, J., A. A. Levchenko, and I. Mejean (2012). Firms, Destinations and Aggregate Fluctuations. RSIE Discussion Paper 627.

Ellison, G., E. L. Glaeser, and W. R. Kerr (2010). What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns. *American Economic Review 100*(3), 1195–1213.

Foerster, A. T., P.-D. G. Sarte, and M. W. Watson (2011). Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production. *Journal of Political Economy 119*(1), 1–38.

Griliches, Z. (1957). Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica 25*(4), pp. 501–522.

Helpman, E. and M. Trajtenberg (1994). A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies. NBER working paper 4854.

Helpman, E. and M. Trajtenberg (1998). *General Purpose Technologies and Economic Growth*, Chapter Diffusion of General Purpose Technologies. Cambridge: MIT Press.

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

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.

Jackson, M. O. and B. W. Rogers (2007). Meeting Strangers and Friends of Friends: How Random Are Social Networks? *American Economic Review 97*(3), 890–915.

Jessen, A. and T. Mikosch (2006). Regularly Varying Functions. *Publications de l'Institut Mathématique 80*(94), 171–192.

Jones, C. I. (2013). Misallocation, Economic Growth, and Input-Output Economics. In *Advances in Economics and Econometrics*, Volume 2. Tenth World Congress: Cambridge University Press.

Jovanovic, B. and P. L. Rousseau (2005). General Purpose Technologies. In P. Aghion and S. Durlauf (Eds.), *Handbook of Economic Growth*, Volume 1, Chapter 18, pp. 1181–1224. Elsevier.

Kelly, B., H. Lustig, and S. Van Nieuwerburgh (2013). Firm Volatility in Granular Networks. Working paper.

Long, J. B. J. and C. I. Plosser (1983). Real Business Cycles. *Journal of Political Economy 91*(1), 39–69.

Murphy, J. B. (1998). Introducing the North American Industry Classification System. *Monthly Labor Review*, 43–47. Technical Note. Bureau of Labor Statistics.

Ngai, L. R. and C. A. Pissarides (2007). Structural Change in a Multisector Model of Growth. *American Economic Review 97*(1), 429–443.

Oberfield, E. (2012). Business Networks, Production Chains and Productivity: A Theory of Input-Output Architecture. mimeo. Federal Reserve Bank of Chicago.

Romer, P. M. (1990). Endogenous Technological Change. *Journal of Political Economy 98*(5), S71–S102.

Vázquez, A. (2003). Growing Networks with Local Rules: Preferential Attachment, Clustering Hierarchy and Degree Correlations. *Physical Review E 67*(056104), 1–17.

Voigtländer, N. (2014). Skill Bias Magnified: Intersectoral Linkages and White-Collar Labor Demand in U.S. Manufacturing. *Review of Economics and Statistics*. forthcoming.

Young, H. P. (2003). The Diffusion of Innovations in Social Networks. In L. E. Blume and S. N. Durlauf (Eds.), *The Economy as an Evolving Complex System*, pp. 267–282. Oxford: Oxford University Press.

APPENDIX

Proof of Proposition 2

Proof. We first derive the probability that the next variety to be classified into any sector s_j sources as an input – indirectly, through its essential inputs – a given individual variety from sector s_i. Recall from Definition 1 that μ_{s_j} is the baseline vector that defines sector s_j. For example, for a car these may be wheels, an engine, and a body. We will refer to an "ideal variety for sector s_j" as a variety that uses exactly the essential inputs in μ_{s_j}. Next, let $i_{s_j}(\leq x)$ be the number of positive entries in vector μ_{s_j} which in turn use variety i as an input.[48] Additionally, let k_{s_j} be the expected overlap between the next variety to be classified into sector s_j and the vector μ_{s_j}, i.e., the expected number of varieties that t has in common with the "ideal variety" for sector s_j. Then the probability that the new variety in sector s_j sources from i via its parents is:

$$p_N \left(k_{s_j} \frac{i_{s_j}}{xm} + \left(m_K - k_{s_j}\right) \frac{d_i^{out}(t)}{t} \right) \frac{m_N}{m_K m} \tag{13}$$

where $m = p_K m_K + p_N m_N$ is the expected indegree for each variety (i.e., the expected number of inputs). Since t draws m_K essential inputs, there are overall $m_K m$ inputs in its network neighborhood. Given that t draws m_N varieties from this network, the term $\frac{m_N}{m_K m}$ gives the probability that it sources any given input via its network of essential inputs. Next, the term in parentheses in (13) gives the probability that a given essential input sources from variety i. This breaks down into two parts. The first term in the parentheses accounts for the possibility that i may be in the network neighborhood of those essential inputs that classify t into sector s_j (i.e., inputs in the set μ_{s_j}). The term gives the probability that t will source from i, conditional on t being classified into sector s_j and sharing, in expectation, k_{s_j} essential inputs with the ideal variety defining sector s_j. The term $\frac{i_{s_j}}{xm}$ gives the probability of drawing i as a network input via these ideal varieties. In expectation, the new variety will have k_{s_j} such draws. The second term accounts for the fact that t may also adopt input i via essential inputs that are not in the set μ_{s_j}, i.e., are not used to classify t as belonging to s_j. This term gives the probability of drawing i as an input via the network, given that $m_K - k_{s_j}$ essential inputs are expected to be drawn uniformly at random from the population. Finally, p_N is the probability that input i is actually adopted by the new variety given that it has been discovered via its essential parents.

Now, according to our definition, each sector is a partition of the set of existent varieties. Hence, the probability that sector s_j starts sourcing from sector s_i at t, conditional on not having done so till $t-1$ is the probability that the new variety t selects as a network input any given variety in

[48] In other words, $i_{s_j}(\leq x)$ is the number of links that lead from the essential varieties defining sector j to variety i.

sector s_i. This is obtained by summing the above expression over all varieties classified in sector s_i:

$$\sum_{i' \in s_i} p_N \left(k_{s_j} \frac{i'_{s_j}}{xm} + \left(m_K - k_{s_j}\right) \frac{d_{i'}^{out}(t)}{t} \right) \frac{m_N}{m_K m}$$
$$= p_N \left(k_{s_j} \sum_{i' \in s_i} \frac{i'_{s_j}}{xm} + \left(m_K - k_{s_j}\right) \sum_{i' \in s_i} \frac{d_{i'}^{out}(t)}{t} \right) \frac{m_N}{m_K m}$$

Finally, note that $k_{s_j} = k$ for all sectors j, i.e., the expected overlap of the new variety t with any sector's 'ideal' list is the same across all sectors. This is immediate from the joint assumption that both ideal varieties defining a sector and the set of essential parents drawn by the new variety are selected uniformly at random from the set of $t-1$ existing varieties. Hence, the expression above simplifies to:

$$p_N \left(k \sum_{i' \in s_i} \frac{i'_{s_j}}{xm} + (m_K - k) \sum_{i' \in s_i} \frac{d_{i'}^{out}(t)}{t} \right) \frac{m_N}{m_K m}$$

For any two sectors, j and j', this expression will only differ in the term $\sum_{i' \in s_i} \frac{i'_{s_j}}{xm}$. Hence if $\sum_{i' \in s_i} i'_{s_j} > \sum_{i' \in s_i} i'_{s_{j'}}$, then j is more likely to adopt a variety in sector i than j'. Now $\sum_{i' \in s_i} i'_{s_j} = \mu'_{s_j} \nu_{s_i} \equiv n_{(s_j, s_i)}$. Thus, if s_j is closer to s_i than $s_{j'}$ at time $t-1$, then j is more likely to adopt from i at time t, as claimed in the proposition. □

Proof of Proposition 3

TBD – see sketch in the text. □

FIGURES

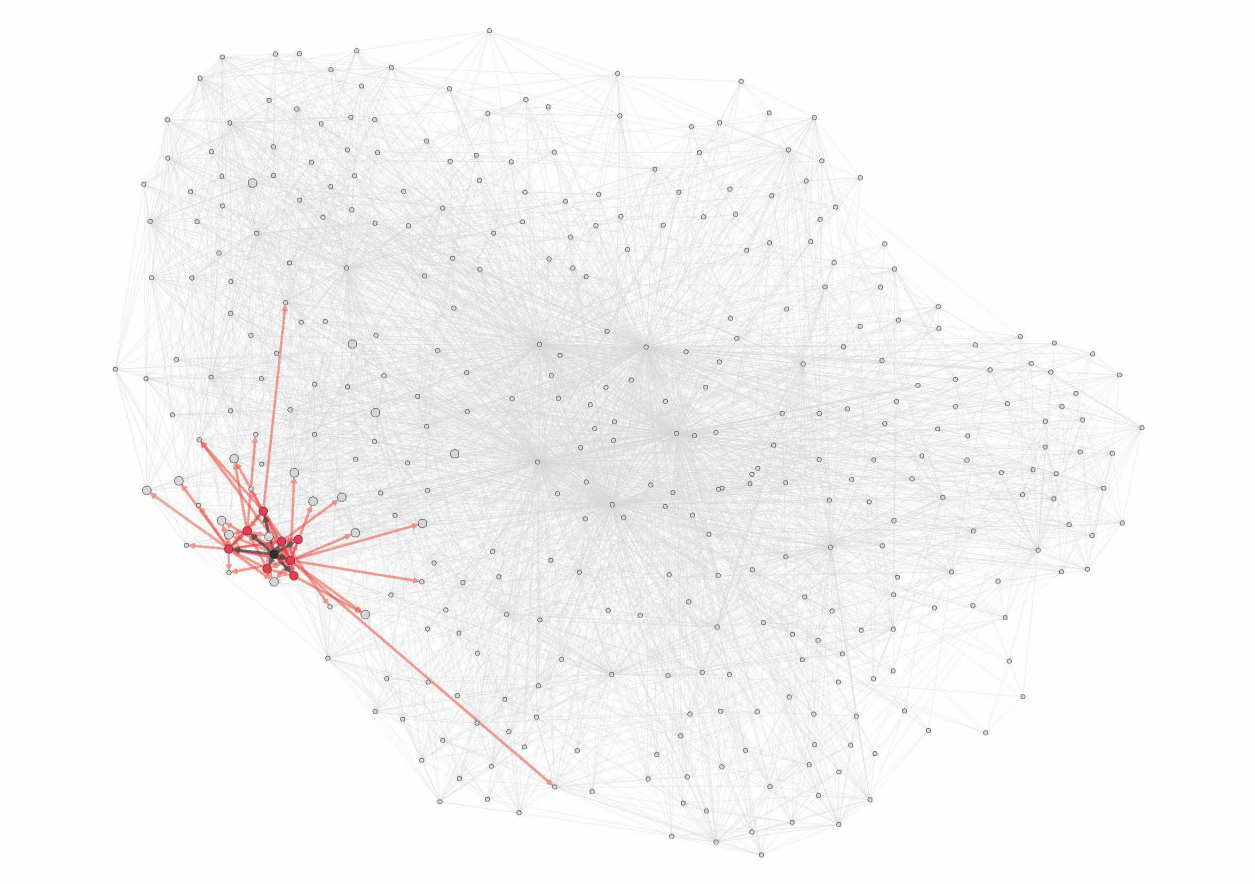


Figure 1: Input-output network and semi-conductor linkages in 1967

Notes: The figure shows the U.S. input-output network in 1967. The black dot represents the semiconductors sector. Red dots are sectors using semiconductors in 1967. Black arrow are flows of semiconductors to using sectors, and red arrows reflect input flows from sectors using semiconductors in 1967 to other sectors.

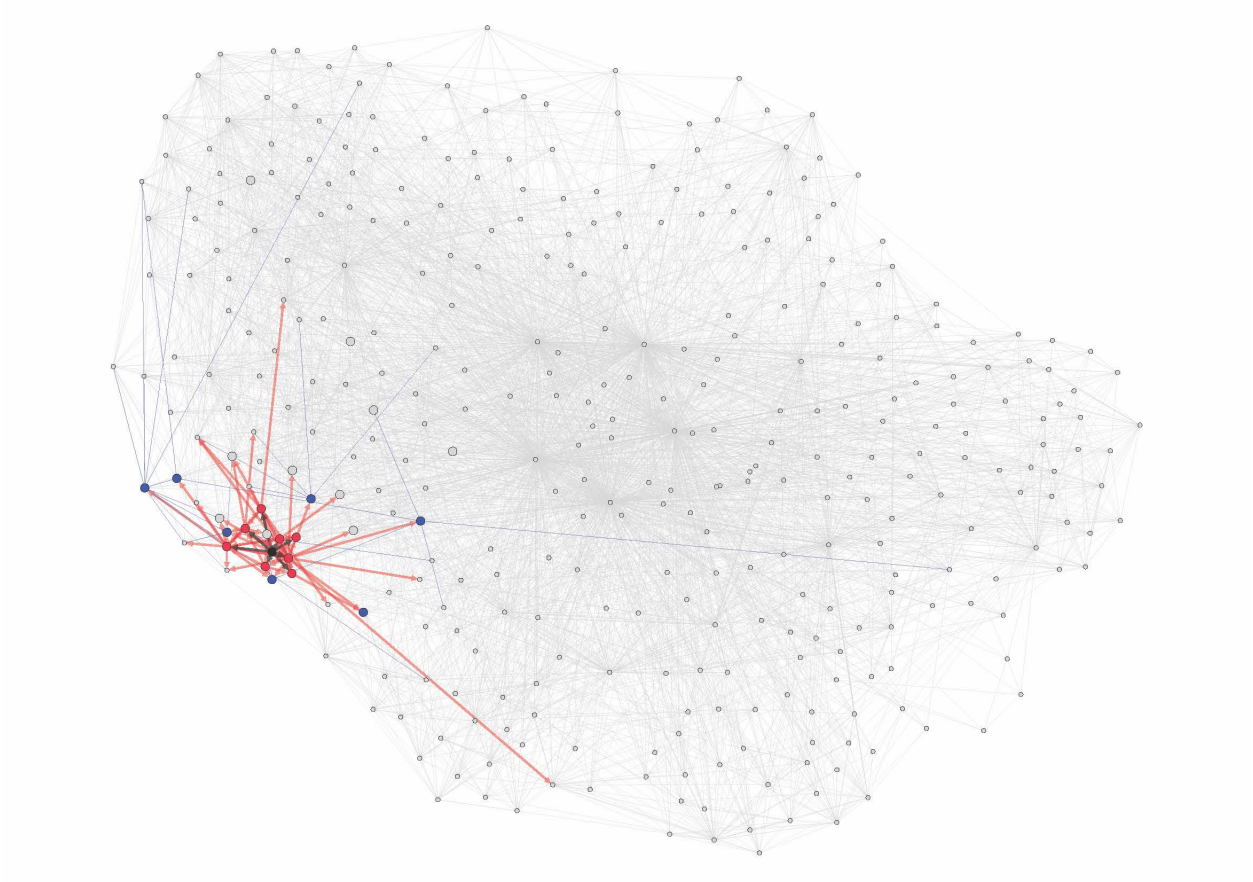


Figure 2: Adoption of semi-conductors in 1972

Notes: The figure shows the U.S. input-output network in 1972. The dots and arrows are the same as in Figure 1. In addition, blue dots represent sectors adopting semiconductors in 1972.

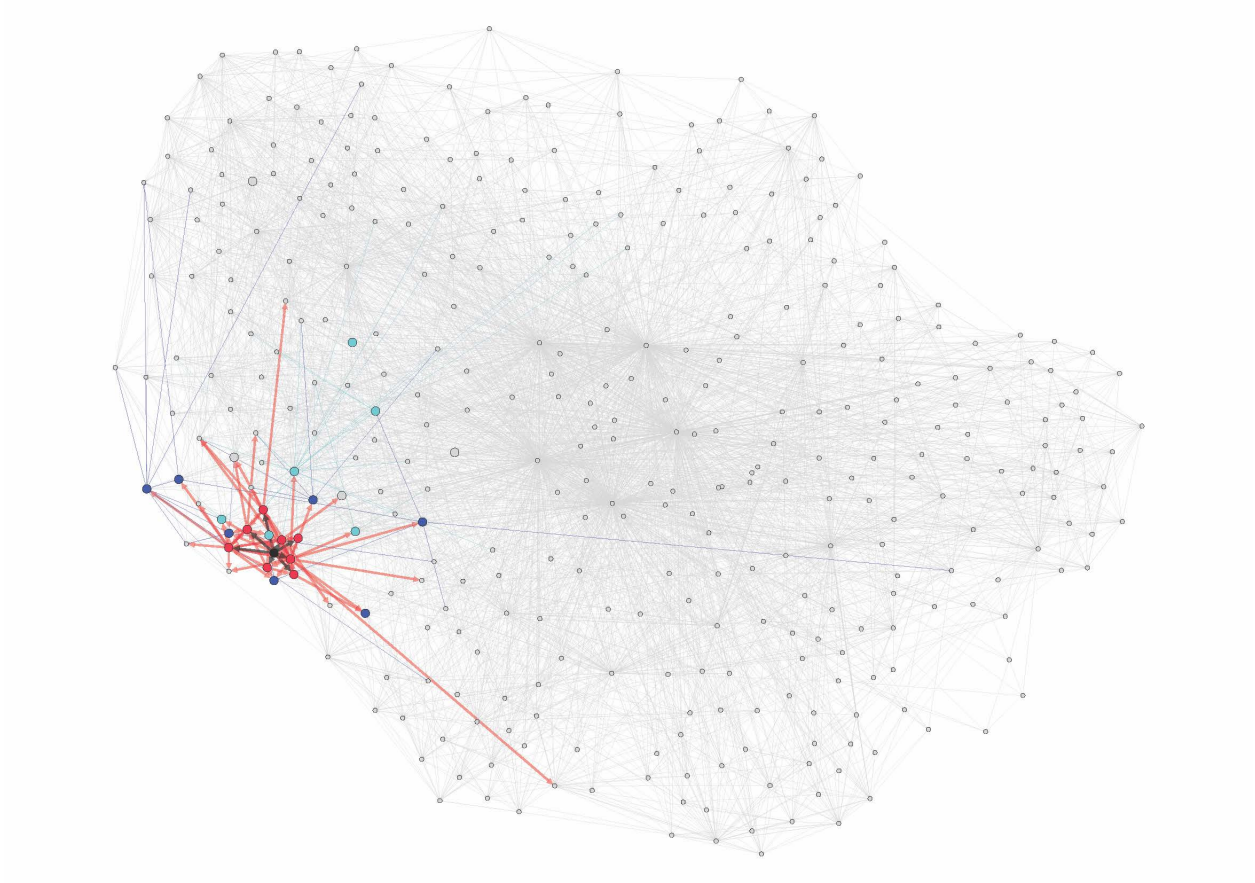


Figure 3: Adoption of semi-conductors in 1977

Notes: The figure shows the U.S. input-output network in 1977. The dots and arrows are the same as in Figure 1. In addition, blue dots represent sectors adopting semiconductors in 1972, and cyan dots are sectors adopting semiconductors in 1977.

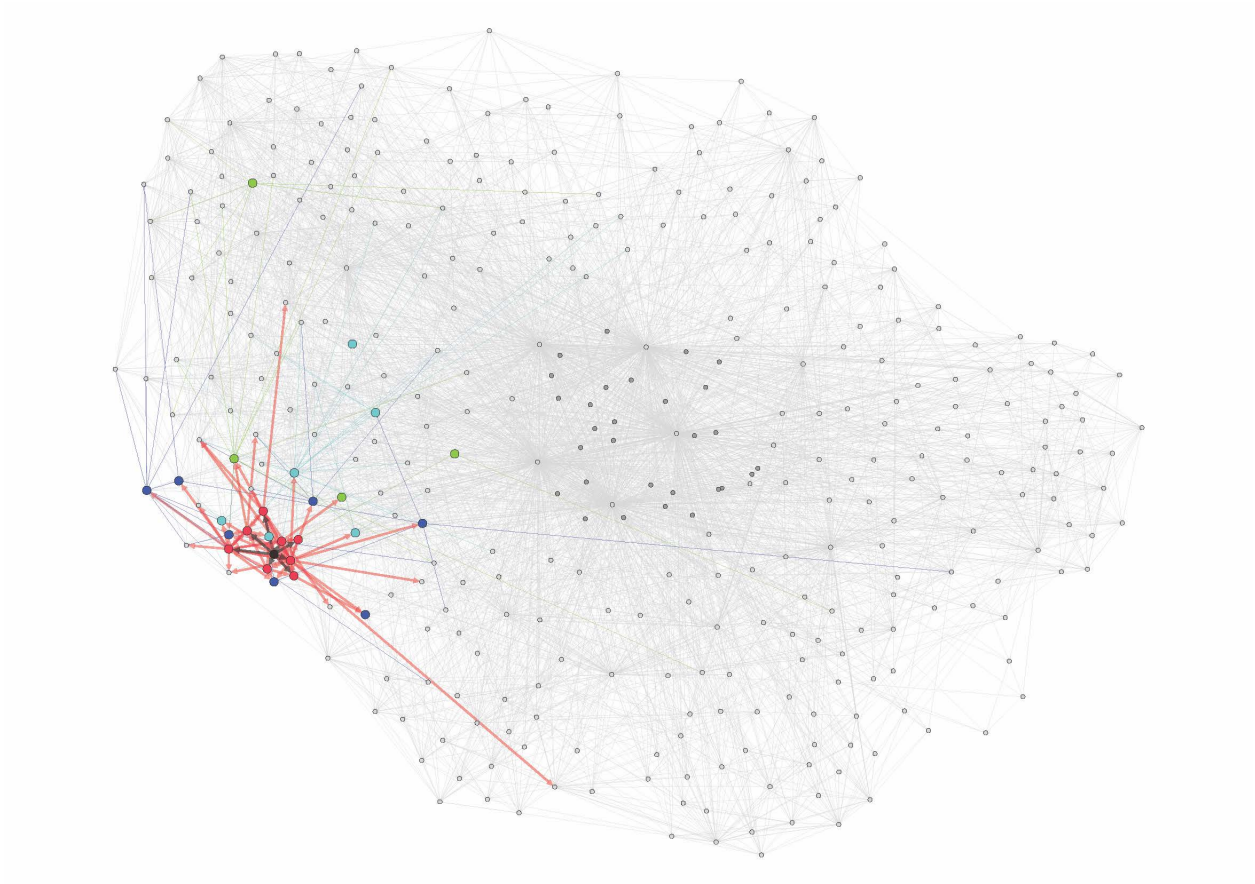


Figure 4: Adoption of semi-conductors in 1982

Notes: The figure shows the U.S. input-output network in 1982. The dots and arrows are the same as in Figure 1. In addition, blue dots represent sectors adopting semiconductors in 1972, cyan dots are sectors adopting semiconductors in 1977, and green dots indicate sectors adopting semiconductors in 1982.

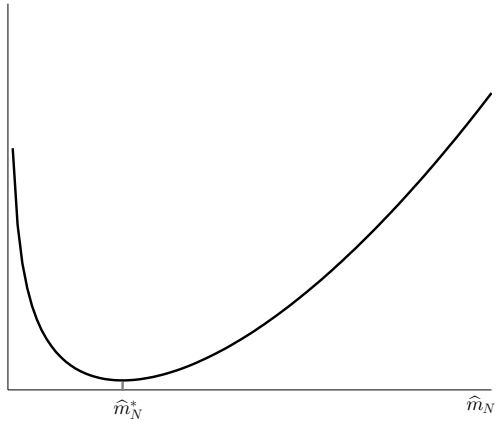


Figure 5: Optimal choice of network input adoption

Notes: The figure illustrates the optimal choice of input adoption. The x-axis shows the number of adopted network inputs, $\widehat{m}_N$. These are ranked by their customization cost as explained in Section 3.2. The y-axis shows the term from equation (8) that is proportional to marginal production cost, and that an input adopter seeks to minimize. For small $\widehat{m}_N$, the input variety effect à la Romer (1990) dominates, so that production costs are decreasing if more inputs are adopted. For higher $\widehat{m}_N$, customization costs for each additional adopted input are also high, outweighing the input variety effect. Thus, production cost become increasing in $\widehat{m}_N$. The optimal number of adopted network inputs is denoted by $\widehat{m}_N^*$.

TABLES

Table 1: Panel on input adoption: Baseline results

Dep. Var.: Dummy for adoption of input i by sector j in year y

	(1)	(2)	(3)	(4)	(5)	(6)
Estimation	Probit	Probit	OLS	OLS	Hazard	Hazard
Distance $d_{ij}(y-5)$	-0.188***	-0.188***	-0.009***	-0.003***	-0..595***	-0.595***
	(0.003)	(0.003)	(0.000)	(0.000)	(0.004)	(0.004)
	[-2.34%]	*[-2.34%]*	*[-1.45%]*	*[-1.45%]*	*[-4.14%]*	*[-4.20%]*
$\triangle_5 TFP_i$		0.113***		0.010***		1.506***
		(0.044)		(0.001)		(0.128)
		[0.07%]		*[0.12%]*		*[0.16%]*
$\triangle_5 TFP_j$		-0.095**		0.005***		1.208**
		(0.042)		(0.001)		(0.0985)
Observations	577,498	577,498	577,498	564,327	577,498	577,498

Notes: The dependent variable is a dummy that takes on value 1 if sector j adopted input i in a given year y between 1972 and 2002. Adoption is defined in Section 4.2; we use the broad definition throughout in this table. The table excludes adoptions occurring in 1997 because of the transition from SIC to NAICS classification in that year. The main explanatory variable is network distance of input i from sector j in the previous I-O benchmark year (i.e., with a 5-year lag), as described in Section 4.1. $\triangle TFP$ denotes the change in total factor productivity over the previous five years in i and j.

Table 2: Additional panel results on input adoption

Dep. Var.: Dummy for adoption of input i by sector j in year y

	(1)	(2)	(3)	(4)	(5)	(6)
Links excluded‡		2-digit			2-digit	
	Broad definition of adoption			Narrow definition of adoption		
Distance $d_{ij}(y-5)$	-0.161***	-0.158***		-0.281***	-0.276***	
	(0.003)	(0.004)		(0.007)	(0.007)	
	[-1.87%]	*[-1.79%]*		*[-1.75%]*	*[-1.68%]*	
Distance in 1967, d_{ij}^{67}			-0.318***			-0.396***
			(0.007)			(0.009)
			[-1.22%]			*[-0.92%]*
TFP_i	0.228***	0.225***	0.013	0.141***	0.131***	0.050*
	(0.021)	(0.021)	(0.023)	(0.025)	(0.026)	(0.028)
	[0.25%]	*[0.24%]*	*[0.01%]*	*[0.08%]*	*[0.07%]*	*[0.03%]*
TFP_j	-0.040*	-0.040*	-0.011	0.011	0.009	0.031
	(0.021)	(0.022)	(0.025)	(0.029)	(0.030)	(0.031)
$\triangle_5 TFP_i$	-0.100**	-0.068	0.129**	1.073***	1.179***	1.315***
	(0.045)	(0.047)	(0.053)	(0.067)	(0.067)	(0.066)
$\triangle_5 TFP_j$	-0.033	-0.037	0.006	0.054	0.087	0.063
	(0.046)	(0.048)	(0.056)	(0.072)	(0.076)	(0.084)
$\ln(emp)_i$	0.123***	0.130***	0.071***	0.176***	0.187***	0.136***
	(0.003)	(0.003)	(0.004)	(0.005)	(0.005)	(0.005)
$\ln(emp)_j$	0.088***	0.092***	0.075***	0.158***	0.163***	0.175***
	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)
Year FE	✓	✓	✓	✓	✓	✓
Observations	570,369	539,235	463,007	478,207	452,139	397,647

Notes: All regressions are estimated by Probit. The dependent variable is a dummy that takes on value 1 if sector j adopts input i in a given year y between 1972 and 2002. Both i and j are observed at the 4-digit SIC level, and the panel extends over the period 1967-2002 in 5-year intervals. Adoption is defined in Section 4.2; columns 1-3 use the broad measure, and columns 4-6 use the narrow measure. The latter requires new i-j links to remain intact for at least 15 years in order to qualify as adoption. The table excludes adoptions occurring in 1997 because of the transition from SIC to NAICS classification in that year. The main explanatory variable is network distance of input i from sector j in the previous I-O benchmark year (i.e., with a 5-year lag), as described in Section 4.1. Columns 3 and 6 use the distance measured in 1967. $\triangle_5 TFP$ denotes the change in total factor productivity in the 5 years prior to each benchmark year (y), and TFP is the level in year y. The number of employees in the sector is denoted by emp. Standard errors in parentheses, clustered at the adopting sector (j) level. * p<0.1, ** p<0.05, *** p<0.01. Values in *[square brackets]* are standardized coefficients, reflecting the change in adoption probability (over a 5-year interval) due to a one standard deviation increase in the explanatory variable.

‡ Columns 2 and 5 exclude all i-j pairs that belong to the same 2-digit industry.

Table 3: Robustness checks – panel estimation

Dep. Var.: Dummy for adoption of input i by sector j; Probit estimation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Years excluded	1997	none	1972,97	1972,97	1997	none	1972,97	1972,97
Links excluded‡				2-digit				2-digit
	Broad definition of adoption				Narrow definition of adoption			
PANEL A: All input relationships								
Distance $d_{ij}(y-5)$	-0.197***	-0.137***	-0.146***	-0.139***	-0.367***	-0.367***	-0.356***	-0.352***
	(0.010)	(0.006)	(0.011)	(0.012)	(0.020)	(0.020)	(0.030)	(0.033)
	[-1.55%]	*[-1.53%]*	*[-1.04%]*	*[-0.95%]*	*[-1.21%]*	*[-1.21%]*	*[-1.03%]*	*[-1.02%]*
TFP_i	0.220***	0.256***	0.352***	0.378***	0.099*	0.099*	0.238***	0.218***
	(0.032)	(0.028)	(0.043)	(0.046)	(0.057)	(0.057)	(0.074)	(0.078)
	[0.19%]	*[0.30%]*	*[0.24%]*	*[0.25%]*	*[0.04%]*	*[0.04%]*	*[0.07%]*	*[0.07%]*
TFP_j	-0.016	-0.087***	0.032	0.009	0.009	0.009	0.017	-0.011
	(0.030)	(0.026)	(0.037)	(0.039)	(0.057)	(0.057)	(0.073)	(0.079)
$\ln(emp)_i$	-0.021	0.005	-0.040**	-0.051**	0.002	0.002	-0.257***	-0.257***
	(0.015)	(0.013)	(0.020)	(0.021)	(0.036)	(0.036)	(0.051)	(0.055)
$\ln(emp)_j$	0.029**	0.038***	0.012	0.027	0.203***	0.203***	0.166***	0.203***
	(0.014)	(0.012)	(0.017)	(0.018)	(0.033)	(0.033)	(0.044)	(0.047)
Using Sector FE	✓	✓	✓	✓	✓	✓	✓	✓
Producing Sector FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	532,473	632,767	436,327	390,061	380,820	380,820	280,898	242,966
PANEL B: Adoption only if at least $1mio input purchase								
Distance $d_{ij}(y-5)$	-0.170***	-0.139***	-0.176***	-0.165***	-0.231***	-0.231***	-0.302***	-0.304***
	(0.010)	(0.009)	(0.013)	(0.013)	(0.021)	(0.021)	(0.029)	(0.032)
	[-0.96%]	*[-0.94%]*	*[-1.03%]*	*[-0.91%]*	*[-0.62%]*	*[-0.62%]*	*[-0.85%]*	*[-0.83%]*
TFP_i	0.132***	0.225***	0.133***	0.148***	0.096	0.096	0.002	-0.059
	(0.039)	(0.036)	(0.048)	(0.052)	(0.060)	(0.060)	(0.076)	(0.081)
	[0.09%]	*[0.17%]*	*[0.08%]*	*[0.08%]*	*[0.03%]*	*[0.03%]*	*[0.00%]*	*[-0.02%]*
Controls as in Panel A	✓	✓	✓	✓	✓	✓	✓	✓
Observations	584,212	681,129	475,629	433,647	398,289	398,289	293,198	253,236

Notes: The dependent variable is a dummy that takes on value 1 if sector j adopts input i in a given year y (in 5-year intervals between 1967 and 2002). Adoption is defined in Section 4.2; columns 1-4 use the broad measure, and columns 5-8 use the narrow measure. The latter requires new i-j pairs to be present for at least 15 years in order to qualify as adoption. Columns 2 and 6 include all benchmark years, including 1997, when the I-O classification switched from SIC to NAICS. The main explanatory variable is network distance of input i from sector j in the previous I-O benchmark year (i.e., with a 5-year lag), as described in Section 4.1. Standard errors in parentheses, clustered at the adopting sector (j) level. * p<0.1, ** p<0.05, *** p<0.01. Values in *[square brackets]* are standardized coefficients, reflecting the change in adoption probability (over a 5-year interval) due to a one standard deviation increase in the explanatory variable.

† Column 5 excludes all i-j pairs that belong to the same 2-digit industry.

‡ Columns 4 and 8 exclude all i-j pairs that belong to the same 2-digit industry.

Table 4: Time to adoption: Baseline results

Dep. Var.: Time to adoption of input i by sector j after 1967

	(1)	(2)	(3)	(4)	(5)	(6)
Years excluded	1997	1997	1972,97	none	1997	1997
Other remarks					2-digit†	narrow‡
Distance d_{ij} in 1967	0.937***	3.112***	1.778***	3.104***	3.222***	1.228***
	(0.196)	(0.341)	(0.360)	(0.311)	(0.354)	(0.290)
	[0.64]	*[2.14]*	*[1.15]*	*[2.04]*	*[2.22]*	*[0.72]*
$\triangle TFP_i(1967 - y_{adopt})$	-96.925***	-364.787***	-331.477***	-281.502***	-363.733***	-146.929***
	(3.919)	(13.186)	(26.434)	(11.861)	(13.695)	(12.575)
	[-1.78]	*[-6.70]*	*[-3.95]*	*[-4.37]*	*[-6.72]*	*[-3.06]*
Using Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
Producing Sector FE	No	Yes	Yes	Yes	Yes	Yes
R^2	0.19	0.73	0.72	0.67	0.73	0.66
Observations	14,849	14,849	8,604	24,312	13,930	6,421

Notes: The dependent variable is the log of years to adoption of input i by sector j after 1967, conditional on this adoption having happened between 1972 and 2002; see equation (10). For a description of network distance d_{ij} see Section 4.1. $\triangle TFP_i(1967 - y_{adopt})$ is the average annual change in TFP in the input producing sector between 1967 and the year of adoption by j. Standard errors in parentheses, clustered at the adopting sector (j) level. * p<0.1, ** p<0.05, *** p<0.01. Values in *[square brackets]* are standardized coefficients, reflecting the change in the dependent variable due to a one standard deviation increase in the explanatory variable.

† Column 5 excludes all i-j pairs that belong to the same 2-digit industry.

‡ The narrow definition of adoption requires new i-j pairs to be present for at least 15 years in order to qualify as adoption.

Table 5: Time to adoption: Additional results

Dep. Var.: Time to adoption of input i by sector j after 1967

	(1)	(2)	(3)	(4)	(5)	(6)
Remarks			2SLS†			narrow‡
Distance d_{ij} in 1967	1.620***	0.968***	0.976***	3.464***	3.148***	0.889***
	(0.181)	(0.212)	(0.182)	(0.323)	(0.327)	(0.229)
	[1.11]	*[0.66]*	*[0.67]*	*[2.37]*	*[2.17]*	*[0.52]*
$\triangle TFP_i(1967 - y_{adopt})$			-211.401***		-147.042***	
			(24.137)		(14.199)	
			[-3.88]		*[-2.70]*	
$\triangle P_i(1967 - y_{adopt})$	99.765***			157.734***	134.913***	130.989***
	(3.029)			(4.282)	(4.760)	(2.740)
	[4.13]			*[6.52]*	*[5.40]*	*[5.31]*
$\triangle TFP_i(1958 - 67)$		-18.341***				
		(6.189)				
		[-0.28]				
Using Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
Producing Sector FE	No	No	No	Yes	Yes	Yes
R^2	0.26	0.17	0.16	0.76	0.77	0.82
Observations	15,072	15,072	14,849	15,072	14,849	6,456

Notes: The dependent variable is the log of years to adoption of input i by sector j after 1967, conditional on this adoption having happened between 1972 and 2002; see equation (10). For a description of network distance d_{ij} see Section 4.1. $\triangle TFP_i(1967 - y_{adopt})$ is the average annual change in TFP in the input producing sector between 1967 and the year of adoption by j; $\triangle P_i(1967 - y_{adopt})$ is the same measure for the price index of i, and $\triangle TFP_i(1958 - 67)$ is the average annual TFP change between 1958 and 1967 in i. Standard errors in parentheses, clustered at the adopting sector (j) level. * p<0.1, ** p<0.05, *** p<0.01. Values in *[square brackets]* are standardized coefficients, reflecting the change in the dependent variable due to a one standard deviation increase in the explanatory variable.

† Two stage least square regression uses historical TFP growth in input-producing sectors ($\triangle TFP_i$ 1958-67) as in instrument for TFP growth after 1967 ($\triangle TFP_i$ since '67). The first stage has an F-statistic of 807.

‡ The narrow definition of adoption requires new i-j pairs to be present for at least 15 years in order to qualify as adoption.