UNIVERSITY OF CAMBRIDGE
Faculty of Economics

Institute for
New Economic Thinking

Cambridge-INET Institute

Cambridge-INET Working Paper Series No: 2015/07

Cambridge Working Paper in Economics: 1550

INPUT DIFFUSION AND THE EVOLUTION OF PRODUCTION NETWORKS

Vasco M. Carvalho (University of Cambridge and CEPR)
Nico Voigtländer (UCLA, NBER, and CEPR)

The adoption and diffusion of inputs in the production network is at the heart of technological progress. What determines which inputs are initially considered and eventually adopted by innovators? We examine the evolution of input linkages from a network perspective, starting from a stylized model of network formation. Producers direct their search for new inputs along vertical linkages, screening the network neighborhood of existing suppliers to identify potentially useful inputs. A subset of these is then adopted, following a tradeoff between the benefits from input variety and the costs of customizing new inputs. Guided by this framework, we document a novel stylized fact at both the sector and the firm level: producers are more likely to adopt inputs that are already used – directly or indirectly – by their current suppliers. In particular, using disaggregated input-output data, we show that initial network proximity of a sector in 1967 significantly increases the likelihood of adoption throughout the subsequent four decades. A one-standard deviation decrease in network distance is associated with an increase in the adoption probability by one third to one half. Similarly, U.S. firms are significantly more likely to develop new input linkages among their suppliers' network neighborhood. Our results imply that the existing production network plays a crucial role in the diffusion of inputs and the evolution of technology.

Input Diffusion and the Evolution of Production Networks*

Vasco M. Carvalho
Cambridge and CEPR

Nico Voigtländer
UCLA, NBER, and CEPR

First Version: March 2014
This Version: February 2015

Abstract

The adoption and diffusion of inputs in the production network is at the heart of technological progress. What determines which inputs are initially considered and eventually adopted by innovators? We examine the evolution of input linkages from a network perspective, starting from a stylized model of network formation. Producers direct their search for new inputs along vertical linkages, screening the network neighborhood of existing suppliers to identify potentially useful inputs. A subset of these is then adopted, following a tradeoff between the benefits from input variety and the costs of customizing new inputs. Guided by this framework, we document a novel stylized fact at both the sector and the firm level: producers are more likely to adopt inputs that are already used – directly or indirectly – by their current suppliers. In particular, using disaggregated input-output data, we show that initial network proximity of a sector in 1967 significantly increases the likelihood of adoption throughout the subsequent four decades. A one-standard deviation decrease in network distance is associated with an increase in the adoption probability by one third to one half. Similarly, U.S. firms are significantly more likely to develop new input linkages among their suppliers' network neighborhood. Our results imply that the existing production network plays a crucial role in the diffusion of inputs and the evolution of technology.

JEL: O33, C67, D57, L23

Keywords: Input adoption, directed network search, dynamics of production networks

*We would like to thank Daron Acemoglu, Ufuk Akcigit, Enghin Atalay, Sanjeev Goyal, Ali Hortaçsu, Ali Jadbabaie, Chad Jones, Ezra Oberfield, Michael Peters, Giacomo Ponzetto, Alireza Tahbaz-Salehi and Jaume Ventura, as well as seminar participants at CREI, the NBER Summer Institute, the NBER EFJK Growth Group, the SED meeting in Seoul, and the University of Cambridge for helpful comments and suggestions. Carvalho acknowledges the financial support of the European Research Council grant #337054, the Cambridge-INET Institute, and the Keynes Fellowship at the University of Cambridge. Voigtländer acknowledges support from the UCLA Anderson Easton Technology Leadership Program. We are grateful to Ali Hortaçsu and Enghin Atalay for sharing their data on supplier-to-customer linkages.

1 Introduction

The adoption of new inputs is an important dimension of technological progress. This is true for both product innovation – where the integration of new inputs leads to new or improved output – and for process innovation, where new inputs can raise the efficiency of production. Input-output linkages are also important for macroeconomic outcomes: they can amplify idiosyncratic sectoral distortions into large aggregate productivity differences (Ciccone, 2002; Jones, 2011, 2013), and they can create aggregate fluctuations by propagating micro-level shocks (Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi, 2012). Yet, input linkages are typically taken as given; little is known about the evolution of production networks.[1]

In this paper, we analyze the formation of input linkages. We ask what determines which inputs are initially considered and eventually adopted in the production of new or improved goods.[2] Similarly, why are some inputs so much more prominent than others?[3] We take a network perspective to answer these questions. To structure our analysis, we build on a standard network formation argument and hypothesize that producers direct their search for new inputs to the network neighbourhood of their existent suppliers. Guided by this stylized model, we explore the empirical determinants of new input link formation both at the sector and the firm level. We uncover a striking empirical regularity: producers are more likely to adopt inputs that are already used upstream – directly or indirectly – by their current suppliers. By the same token, we find that inputs that are initially closer to many potential adopters are more likely to diffuse widely. Our results imply that the existing input-output network plays a key role in the formation of new linkages.

To guide our analysis, we provide a stylized model of network formation at the variety level and then explore its sector-level implications. Each variety producer is embedded in a network of production linkages – producers do not only interact directly with their suppliers, but also indirectly with input producers further upstream. This 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) directly used. In order to keep the analysis simple, we take the arrival of new varieties as given and focus on their input adoption decisions. In each period, a new variety emerges exogenously. It then

[1] The exceptions are Atalay, Hortaçsu, Roberts, and Syverson (2011) and Oberfield (2013); both examine the evolution of linkages at the firm level. For a recent overview of the literature on production networks see Carvalho (2014).

[2] Firms often experiment with several potentially suitable inputs before making their final choice. For example, Steve Jobs famously had the first iPhone's screen changed from plastic to hardened glass only four weeks before mass production began in 2007.

[3] The number of sectors that source inputs from a given supplier follows a power law (Carvalho, 2010). Atalay et al. (2011) and Kelly, Lustig, and Van Nieuwerburgh (2013) report evidence on the distribution of supply linkages at the firm level. The power law in the outdegree distribution is crucial for linkages to augment idiosyncratic shocks into aggregate fluctuations (Acemoglu et al., 2012).

forms input linkages following three steps, where the first two build on the central mechanism of a class of dynamic network formation models. We illustrate these steps in the graph below. First, a new variety producer (j) draws a set of 'essential' input suppliers (nodes g and h). Second, in order to customize its new variety, j searches for further potentially useful inputs among the suppliers of g and h (i.e., among nodes a-e). In other words, the search is directed vertically in the production network, towards the technological neighborhood of essential inputs. Third, the new variety producer j decides which inputs to adopt among those identified in the second step. This decision is driven by a trade-off between the 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 (indicated by the dashed arrows).

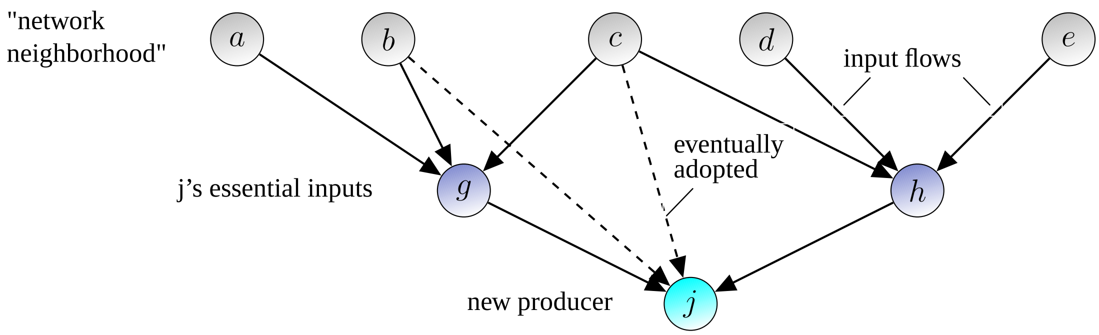

This process implies that individual producers are more likely to adopt inputs that are closer in their network neighborhood. This is similar to the evolution of social networks, where new friendships are more likely to form with friends of friends than with random people. In the context of production, a firm is more likely to adopt an input that its supplier is already using, than a random input from the product universe. The formation of new linkages delivers a law of motion, where the current production network and its evolution are closely interrelated: on the one hand, present network distances determine input adoption; input adoption, on the other hand, changes network distances. This gives rise to a dynamic evolution of the input-output network. The stylized model also delivers a power law in the number of varieties supplied, in line with the empirical regularities observed by Carvalho, Nirei, Saito, and Tahbaz-Salehi (2014).

We then explore the sector-level implications of the variety-level 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 will be assigned to the motor vehicles sector. We show that, based on this definition, the model predicts i) new input linkages across sectors are more likely to emerge within the proximity of existent input supply relations and ii) the power law distribution of forward linkages aggregates up to the sector-level. 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 model's predictions.

Following the theoretical framework, our empirical analysis examines the determinants of input adoption at the sector and firm level. We first 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. We find that sectors are substantially more likely to adopt inputs that are initially closer in their input-output network. This is illustrated in Figure 1. On the x-axis we plot the simplest possible measure of network distance in 1967 – the smallest number of directed input links separating a (potential) input supplier i from a (potential) input adopter j. This provides a simple metric for the vertical distance between an upstream supplier and a downstream potential user of the input. For example, in 1967, the sector "Primary Batteries, Dry and Wet" (SIC 3692) had distance 2 to the (potential adopter) sector "X-Ray Apparatus" (SIC 3844), while it had distance 4 to "Cigarettes" (SIC 2111). The y-axis gives the frequency of input adoption events (j adopting i) observed after 1967.[4] For example, in our data we observe that "X-Ray Apparatus" producers adopted "Primary Batteries, Dry and Wet" as an input in 1977, while "Cigarettes" never did so. The pattern in Figure 1 is striking. Input adoption is much more frequent for sector-pairs that were already relatively close (but not yet directly trading inputs) in the 1967 input-output network. For 22% of sector-pairs that were two input-links apart in 1967 (distance 2), j adopted i over the subsequent 35 year period.[5] This is more than double the frequency of adoption observed for distance-3-sectors (9%), and more than 5 times the frequency observed for sectors that were 5 links apart.

Our main 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 alternative definitions of network distance and adoption. They are also unaffected by a host of controls such as size, proxies for technological progress, as well as fixed effects for adopting and input-producing sectors. Throughout, we document economically significant magnitudes; for example, a one-standard deviation decrease in network distance raises the adoption probability in any given 5-year period by one third to one half.

[4] We say that a sector j adopts input i if there was no input flow from i to j in 1967, and such a flow is recorded at any point thereafter in our sample, which extends until 2002.

[5] Not all of these 'adoptions' are long-lived; many reflect one-time input flows between sectors. All our findings are robust to alternative definition of adoption. In particular, our most conservative definition requires input flows of at least \$1million over at least 15 years. In this case, the frequency of adoption for distance 2 is 4.9%, and – consistently with the differential pattern in Figure 1 – it is significantly lower for larger distances.

Next, we examine the formation of supplier linkages at the *firm* level. We use data from Compustat, which includes information on major customers – those that are responsible for more than 10% of a given seller's revenues. Because of this reporting threshold, the analysis is naturally limited to relatively important links. Based on this data, we construct a network of suppliers and their customers. We confirm that firms are significantly more likely to adopt inputs that have previously been used by their suppliers (close network proximity) than inputs from outside their network neighborhood. Additionally, we show that our findings are robust to the inclusion of firm-level controls, such as fixed effects, firm size, labor productivity, or the geographical distance between firm pairs.

Our analysis shares a common limitation with other studies of production networks: exogenous variation for input-output linkages is not available. This raises the concern of omitted variable bias. For example, a general trend towards computerization may be accompanied by a gradual spread of electronic components as inputs in production. Since these are in turn connected to semiconductors, this process would bring sectors closer to the latter in the input-output network, with some of them eventually directly adopting semiconductors. While this would confound our panel results, it is less likely to affect our results that are based exclusively on initial network distance in 1967. In fact, technological trends would only affect these results if they were related to initial network distance. But this, in turn, is the core of our argument – initial network distance matters for the future evolution of linkages. In addition, we show that our results also hold when we include sector-pair fixed effects, so that we exploit only *changes* in network distance, i.e., the variation that is due to new links forming or existing links disappearing over time. Consequently, unobserved correlates of *initial* network distance are also unlikely to explain the pattern in the data. These findings impose restriction on interpretations of our results: candidates to explain the empirical regularities have to be correlated with network distance (both in levels and changes), and be related to direct adoption of inputs. We discuss three interpretations that fit this pattern.

First, network distance may reflect technological distance in the sense that production processes are more or less similar, rendering 'closer' inputs more compatible. For example, the production of valves is technologically closer to vehicles than food processing, making the former a more feasible input in car production. Second, network distance could proxy for spatial distance to the extent that industries and firms that trade inputs intensively tend to coagglomerate (Ellison, Glaeser, and Kerr, 2010). Third, network proximity may reflect more frequent social interactions through which information about potentially useful inputs is transmitted. A second limitation of our analysis is that, ultimately, we cannot distinguish between these mechanisms. Nevertheless, we can narrow down possible interpretations. The fact that geographical distance between firms does not change our results makes coagglomeration unlikely as a main driver. At the sector level, we show that excluding linkages formed within the same 2-digit sector does not affect our findings; we also show that the forward-distance between sectors (i.e., links via buyers, instead of suppliers) does not predict input adoption. This implies that horizontal similarity of sectors is not a likely candidate to explain the pattern in the data. Rather, our results suggest that *vertical* distance along supply chains is a useful starting point to understand patterns of input adoption and diffusion.

We build on a rich research agenda that has studied the diffusion of technology, starting from the seminal work by Griliches (1957). 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.[6] Our results imply that occupying a relatively central position in the production network – for example, when a new input is used by already central technologies – makes wide diffusion 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.[7]

We also naturally relate to a literature that models the evolution of technology as a recombinatoric process of existing ideas into new ones (Weitzman, 1998).[8] The large number of possible combinations led Weitzman to the conclusion that "the ultimate limits to growth may lie not so much in our abilities to generate new ideas, as in our abilities to process to fruition an ever increasing abundance of potentially fruitful ideas" [p.359]. This begs the question of how innovators organize their ‘search process’ among the myriad of possible combinations. Our approach makes

[6]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.

[7]As discussed above, proximity in the input-output network may also reflect more 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. See Fafchamps, Goyal, and van der Leij (2010) for a study of the determinants of co-authorship patterns in economics, which uses an empirical strategy akin to ours.

[8]Ghiglino (2012) emphasizes that the quality of ‘parental’ ideas plays an important role in this setting. On a related point, Acemoglu, Akcigit, and Kerr (2014) show that downstream technologies that cite upstream technologies with rapid patent growth in the past, are themselves more likely to exhibit subsequently faster innovation.

the object of this search tangible: we view technology as 'production recipes' that prescribe a combination of different inputs to produce output. Correspondingly, the search for "ideas" reflects the combination of existing inputs into new products. Additionally, we provide evidence that this process of recombination of inputs does not occur at random. Rather, it is directed towards inputs that are relatively close in the production network.

Our focus on input-output networks is also related to a literature that emphasizes the role of intersectoral linkages in determining macro-economic outcomes, such as productivity and aggregate fluctuations (Jones, 2013; Carvalho, 2010; Acemoglu et al., 2012; Bigio and La'O, 2013).[9] Further afield, input-output linkages also have important implications for the organization of production and the optimal allocation of ownership rights along global supply chains (Antràs and Chor, 2013; Costinot, Vogel, and Wang, 2013). These literatures invariably take the input-output network as an exogenously given restriction on production technologies, while we examine its evolution.

Our work also builds on a literature of dynamic network formation models (Vázquez, 2003; Jackson and Rogers, 2007; Chaney, 2014). 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. Finally, our paper is closely related to Atalay et al. (2011) and Oberfield (2013), who model input link formation in buyer-supplier networks. Atalay et al. (2011) estimate a model where new links form in part randomly, and in part due to preferential attachment (to prominent, but not necessarily nearby suppliers). Oberfield (2013) provides a mechanism whereby producers randomly search for the lowest cost input supplier. In these mechanisms, firms do not exploit existing supply linkages to search for new inputs; in contrast, we emphasize the role of linkages in *directing* the search for potential inputs. Relative to the existing literature, we are the first to document the novel stylized fact that input adoption is strongly associated with proximity in the production network – and that, consequently, the existing production network plays an important role in its subsequent evolution.

The paper is organized as follows. Section 2 uses the diffusion of semiconductors as a case study of input adoption in a network. Section 3 describes our model of input adoption, starting at the product variety level and then aggregating these into sectors. Section 4 introduces our measure of network distance and describes our sector-level data. In Sections 5 and 6 we present empirical results at the sector- and the firm-level, respectively. Section 7 concludes.

[9]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 (2014) for empirical evidence supporting these mechanisms.

2 Input Diffusion in a Network: The Case of Semiconductors

The diffusion of semiconductors, a key general purpose input, provides a telling illustration of input adoption in a network. Figure 2 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 the sector "Semiconductors and related devices". The red nodes mark sectors that directly sourced semiconductors as an input in 1967 – only a handful of technologies did so. Finally, the red arrows point to indirect users of semiconductors with distance 2, i.e., sectors that sourced inputs which in turn used semiconductors.

Given this starting point, Figures 3-5 show the path of diffusion of semiconductors across sectors over the subsequent 15 years.[10] Blue dots in Figure 3 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 3. Cyan and green dots in Figures 4 and 5 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 subsequent direct adoption of semiconductors as an input.

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, early adoption of semiconductors was strongly associated with initial network proximity.

It is instructive to focus on one of these paths of adoption to illustrate the role of linkages across sectors in the diffusion of semiconductors. In 1967, the "Computers and Office Equipment" sector did not yet directly source semiconductors as an input. However, computers used "Electronic Components", which in turn used "Semiconductors and Related Devices." That is, computers were distance 2 from semiconductors. In the early 1970s, new computer varieties increasingly used the newly developed integrated circuits (a good classified in "Electronic Components", which in turn used semiconducting materials intensively). The increasing reliance on integrated circuits

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

was accompanied by a direct use of semiconducting materials in computer production.[11] Correspondingly, the "Computers and Office Equipment" sector adopted semiconductors in 1972 in our data.[12]

3 A Model of Input Diffusion in a Network of Technologies

In this section we present a stylized 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 (2014), we begin by describing the set of feasible inputs available to each new variety. 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 first draws 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.

[11] 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.

[12] This argument can be extended further – to a sector that had distance 3 from semiconductors in 1967: the "Scales and Balances" sector sourced early computer varieties (but not semiconductors) in the late 1960s to store and perform calculations on weighting 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 production of industrial and retail digital scales. These eventually incorporated semiconductors directly by 1977.

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 $b_{ij} \in \{0, 1\}$ indicates 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.[13]

This production network evolves over time as new varieties arrive sequentially in the economy. In particular, at each time t a new variety is added to the economy.[14] 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. For example, a car producer may consider several different engine options.

In a second step, the new variety can adopt further inputs. This reflects a stage of refinement of variety t by adding features beyond the essential ones. To identify potentially useful inputs, the producer directs its search to inputs that are already used by its essential suppliers. In the car example, the producer may look for options to make the body lighter (e.g., by using ultra-light carbon fiber) or add luxury features to the car. This second round search delivers a spectrum of potential inputs, and only a subset of these will eventually be adopted. To formalize the process of networked 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 receive *information* 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.[15]

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

[14] 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' depending on the context.

[15] The outdegree of i gives the number of varieties to which i supplies, i.e., the number of other varieties that use

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:

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

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.[16] 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.[17] 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 that a given input i is drawn. Note that the same input i can show up several times in t's network neighborhood – via different essential inputs. In our car example, 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.[18]

---

variety i as an input. In contrast, the indegree of i is the number of inputs that i itself uses.

[16]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.

[17]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.

[18]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

Altogether, the second term in equation (1) thus captures the odds of i being adopted by the new variety t through the latter's network search. 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 the 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.[19] 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

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.

[19] 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. In fact, following Dorogovtsev, Mendes, and Samukhin (2000, equ. 9) it is possible to derive a closed-form solution for $F(d^{out})$ without appealing to a mean field approximation. Based on these expressions, Dorogovtsev et al. (2000, equ. 11) show that in the limit of large d^{out}, the distribution implied by the mean-field approximation above is indeed correct. We thank Enghin Atalay for bringing this point to our attention.

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 $\widehat{K}_t \subseteq K_t$ 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.[20] 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 $k \in K_t$, so that $\widehat{K}_t = K_t$.[21]

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

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

[21] 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 ($\widehat{K}_t \bar{x}_{Kt}$).

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.[22] This ensures that our results are not driven by scale effects.[23] We can now specify 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.[24] 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 akin to a tax on output, used to cover the initial adoption cost.[25]

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[26]

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

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

[23] 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 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.

[24] 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.

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

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

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 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 examine the symmetric case, we assume that each variety has the same technology draw $A_t = A$, and that final demand is such that the price of each variety is a constant markup over its marginal cost.[27] In addition, recall that 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. However, we can 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

[27] 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 $\epsilon/(\epsilon-1)$.

expected 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 in prices and indegree is compatible with asymmetry in the number of forward linkages (i.e., the outdegree).

Under symmetry of prices ($\phi_n = \phi,\ \forall n$), and given the above ranking of customization costs, (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 $\beta < 2(\epsilon - 1)$.[28] This delivers a U-shape with a unique minimum (see Figure A.1 in the appendix). 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 eventually 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).

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, we can show that variety-level predictions hold at the sector level.

[28]For example, suppose that the overall expenditure share for intermediate inputs is 0.5, and that half of these are network inputs. This implies $\beta = 0.25$. 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.

3.3 Sector Level Implications

Our model of networked input adoption is defined at the variety level. However, the most prominent source for production network data is available at the sector level in the form of input-output tables. In order to render these data useable for our purposes, we now explore the sectoral implications of our model.

Aggregation of Varieties into Sectors

We start by defining how varieties are assigned to sectors in the context of our model, employing a principle of similarity of inputs. As a result, sectors are composed of varieties that share similar input bundles. This input-based approach is also a guiding principle of actual sectoral classification systems like NAICS.[29] 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 'body', 'engine', and 'wheels'. In the context of our model, the vector μ_{s_j} represents 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.[30] 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 that this definition allows for overlap among sectors, in that different sectors can share some elements across their baseline vectors. For example, 'wheels' may be represented in both the bicycle and car baseline vectors. 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

[29] 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).

[30] 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.

varieties that have been classified into sector s_j and source inputs from varieties classified into sector s_i.[31]

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 $a_{s_is_j} = 0$ at time t. The network proximity of (s_i, s_j) is defined as $n_{(s_i,s_j)} \equiv \mu'_{s_j}\nu_{s_i}$ where ν_{s_i} is a $t \times 1$ vector, with each entry $\nu_{s_i}(v)$ giving the number of varieties from sector s_i that are sourced as inputs by variety v, for $v = 1, .., t$. Then sector s_i is closer to s_j than to $s_{j'}$ if $n_{(s_i,s_j)} > n_{(s_i,s_{j'})}$.*

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. For example, if s_j is the "vehicles" sector then its defining varieties will include body parts. If body parts, in turn, source many steel varieties (from sector s_i="steel") then this will imply a relatively high proximity of "steel" to "vehicles". Formally, $n_{(s_i,s_j)}$ 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.[32] 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 $n_{(s_i,s_j)}$ 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 sector s_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 $n_{(s_i,s_j)} \equiv \mu'_{s_j}\nu_{s_i}$ 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.

[31] For a fixed number of sectors J, as t becomes large, eventually all sector pairs will exhibit non-zero flows $a_{s_is_j}$. We study sector-level adoption, meaning that $a_{s_is_j}$ goes from zero to positive. We thus implicitly assume that the time t input-output network is sparse, i.e., that many $a_{s_is_j}$'s are zero. In addition, note that economies with zero $a_{s_is_j}$ can be maintained even for large t if the sectoral classification system is expanded by raising J, i.e., by refining the sectoral detail.

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

Proposition 2. *Take any two sectors s_j and $s_{j'}$ that previously did not source inputs from sector s_i, i.e., $a_{s_i s_j} = a_{s_i s'_j} = 0$ at $t-1$. If at time $t-1$ sector s_i is closer to s_j (i.e., $n_{(s_i,s_j)} > n_{(s_i,s_{j'})}$), then 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 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.[33] 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). Note that the induced sector 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. This, in turn, 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. *In the symmetric equilibrium, if the variety-level outdegree distribution at time t is power law, so is the distribution of sectoral weighted outdegrees.*

[33]More generally, the new variety t can form links to inputs in sector s_i directly – by drawing variety $i \in s_i$ 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.

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 build on the fact that at any time t, the number of varieties in each sector follows a Binomial distribution. Second, we note that under the assumption of price symmetry, the sectoral (weighted) outdegree is proportional to the total number of varieties to which a sector supplies inputs at time t, where the constant of proportionality is given by the price ϕ. 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) random sum of power law distributed variables. Based on known results on the behavior of random sums of power law variables (see Jessen and Mikosch, 2006, Lemma 3.1) we can then show that any sector's weighted outdegree is itself power law distributed with the same tail exponent as the variety-level outdegree distribution.

4 Empirical Framework and Sector-Level Data

While the core mechanism of the model works at the variety level, our aggregation results in Proposition 3 allow us to employ sector level data. Because the most reliable data are available at the sectoral level, these are the basis for our main empirical analysis. 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. We also provide falsification tests such as network distance following *forward*- (as opposed to backward) linkages. To save on notation, we use j (instead of s_j) to denote the input-using sector, and i (instead of s_i) for the input-producing sector.

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.[34] Crucially,

[34] 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

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 inputs from sector i in the total intermediate input expenditures of sector j. If Γ_{ij} is non-zero, we define the distance from i to j 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 itself, 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 $d_{ij} = d_{ik} + d_{kj}$.[35] Formally, for two sectors i and j that are not directly connected, the shortest path is given by:

$$d_{ij} = \min_{k \neq i} \left\{ \frac{1}{\Gamma_{ik}} + d_{kj} \right\} \tag{9}$$

As this equation shows, 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) 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.

measure of distance captures both these features.

[35] 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 by following forward linkages from sector j to sector i, d_{ji}. Our use of d_{ij} reflects the notion of distance implied by our model, which focuses on network proximity via linkages to input providers (upstream), as opposed to customers (downstream).

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.[36] 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.[37] 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.[38] To avoid that this affects our results, we exclude new linkages formed in 1997 in our baseline analysis. Nevertheless, in the robustness analysis we show that our results hold 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:

$$A_{ij}(y) = \begin{cases} 1, & \text{if } \Gamma_{ij}(y) > 0 \text{ and } \Gamma_{ij}(y') = 0, \; \forall y' < y \\ 0, & \text{otherwise} \end{cases} \tag{10}$$

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

[36]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 as one sector, 'paper and paperboard mills' over the full sample period.

[37]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.

[38]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.

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.[39] 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:

$$T_{ij} = y_{Adopt} - 1967 \,, \tag{11}$$

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 19,885 adoptions in our broad measure and 8,765 in the narrow one.[40]

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).[41] We use these data to derive control variables for input producing and adopting sectors. We also calculate changes in TFP for input producing sectors, $\triangle TFP_i$ 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, $\triangle TFP_i^{58-67}$, strongly predicts $\triangle TFP_i$ after 1967.

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

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

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

5 Sector-Level Evidence

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 sector j adopting input i by 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 *conditional* on adoption occurring, it tends to happen earlier for smaller initial network distance d_{ij}^{67}. This analysis includes only sector pairs for which adoption occurred over our sample period. It thus addresses the potential concern that our results may be driven by the *absence* of adoption events for technologically very distant sectors (such as vehicles and processed food). Instead, our time-to-adopt results – by exploiting only variation among actual adoptions – suggests that the network distance to *feasible* potential inputs plays an important role. We also show that, in line with our model, more rapid technological progress in an input producing sector goes hand-in-hand with higher 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 $d_{ij}(y-5)$:

$$Prob\,(A_{ij}(y)=1) = g\,(\ln d_{ij}(y-5), X_i(y), X_j(y)) \ , \tag{12}$$

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 (10).[42] 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 hazard models, 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 our

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

results, 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 -2.31 percentage points, the effect of network distance is economically significant.[43] 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 0.04 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 (at least over the short 5-year horizon that we analyze here).[44] Finally, there is no clear relationship between TFP growth of adopting sectors ($\triangle TFP_j$) and input adoption.

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 in $d_{ij}(y-5)$ is associated with an increase in the probability of adoption by 1.41 percentage points throughout the following five years. 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.

In columns 5 and 6 we estimate a proportional hazard model. The hazard ratio for distance implies that as $d_{ij}(y-5)$ increases by one unit, the rate of adoption in any given period will be 0.594 as high as before, i.e., it will decrease by 40.6%. Alternatively, a one std increase in $d_{ij}(y-5)$ reduces the adoption rate by 56.3%.[45] The corresponding standardized *relative* hazard coefficient is -4.12 percentage points, implying that over the entire sample period, a one standard deviation increase in network distance is associated with a 4.12 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. The magnitude of both TFP effects remains small, with standardized coefficients in the range of 0.1%. In sum, the hazard model confirms the economically and statistically significant (negative) relationship between initial network distance and the odds of input adoption, as well as the quantitatively small positive effect of TFP growth in

[43]The marginal effect implied by the Probit coefficient in column 1 is -0.0145, and the standard deviation of network distance is 1.59.

[44]Another possible explanation for the small effect is that input producers may charge strategically low prices in order to attract customers (which may be particularly true during the 5 years leading up to adoption). Because our data measure changes in *revenue* TFP, they would understate actual (physical) efficiency growth in the presence of falling prices (Foster, Haltiwanger, and Syverson, 2008; Garcia-Marin and Voigtländer, 2013).

[45]The coefficient is ln(0.594)=-0.521, and thus exp(-0.521*1.59)=0.437 is the hazard ratio for a one std (1.59) increase in network distance.

input producing sectors.

Additional controls, robustness, and forward linkage distance

In Table 2 we present alternative specifications for our probit regression, including (benchmark) year fixed effects and additional controls. Columns 1-4 use our broad measure of input adoption; columns 5-8 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 6 exclude input links that are formed between 4-digit sectors within the same 2-digit industry. This reduces the number of adoption events by 9%.[46] Thus, most input adoptions occur *across* 2-digit sectors. Second, columns 3 and 7 use network distance measured at the beginning of the sample period, in 1967, d_{ij}^{67}. All regressions in Table 2 now 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.

We find that neither the additional controls nor the variations in the network distance measure change our results. Throughout all specifications, network distance is strongly negatively associated with 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 1.8 percentage points. Note that the coefficients on $d_{ij}(y-5)$ are almost unchanged when we exclude linkages within 2-digit industries (columns 2 and 6). This makes it unlikely that horizontal similarity of sectors is driving our results. When using distance in 1967 (columns 3 and 7), a one std reduction in d_{ij}^{67} (0.65) raises the probability of adoption by approximately 1.2 percentage points. This 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. Importantly, the fact that our results remain strong when we use d_{ij}^{67} suggests that unobserved trends are unlikely to be a major confounding factors. When including the time-varying distance $d_{ij}(y-5)$ together with the initial distance d_{ij}^{67} (columns 4 and 8), we find that both are significantly positively associated with the probability of adoption.

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 0.25 percentage points.[47] On the other hand, the coefficients on efficiency of the

[46]For the broad (narrow) measure, we count 19,885 (8,765) input adoption events in our sample (excluding 1997), and this number declines to 18,111 (8,054) when excluding adoption events within 2-digit industries.

[47]We thus confirm our previous finding that network proximity is the quantitatively 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 1.74 (0.08) 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

adopting sector have ambiguous signs and are mostly insignificant. Finally, sector size (measured by employment) is associated with both higher probability of adopting and being adopted.

Is the observed relationship between network distance and input adoption 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.[48] Both significance and magnitude of the coefficient on network distance are unchanged. Remarkably, the standardized coefficient of initial network distance d_{ij}^{67} is now larger than the one for the time-varying measure $d_{ij}(y-5)$ (columns 4 and 8). This implies that, once sectoral idiosyncracies are filtered out, distance in 1967 is strongly associated with input adoption even as the network itself evolves over time.[49] In other words, in line with our argument, the *initial* network structure of the production network provides strong predictive power for its long-run evolution. The evidence on TFP in input producing sectors (TFP_i) is now mixed, with mostly positive but quantitatively small coefficients. Among the controls that are not separately reported in Table 3, TFP_j (for adopting sectors) shows no clear relationship with the likelihood of input adoption, and the relationship between input adoption and employment is now ambiguous for input producing sectors (i), and less robust than above for adopting sectors (j). Panel B of Table 3 documents very similar results when we restrict adoption events to purchases above $1 million. This ensures that minor transactions in input-output tables do not drive our results. Note that in our most restrictive specification (using the narrow definition of adoption) in column 8, only network distance d_{ij}^{67} is strongly negatively associated with input adoption, while $d_{ij}(y-5)$ is insignificant and positive. This provides further support for our focus on initial network distance.

In Table 4 we include more restrictive *pairwise* fixed effects for each i-j combination. Our baseline analysis does not include these because we emphasize the role of *initial* network distance in explaining the subsequent evolution of linkages, and i-j fixed effects effectively filter out the distance in 1967. However, including i-j effects also offers an advantage: it allows us to restrict the identifying variation to *changes* in network distance. We can thus examine whether shortened network distances due to previous input diffusion raised the subsequent likelihood of adoption. For example, semiconductors were 2 nodes away from "scales and balances" in 1967, with 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.).

[48]The incidental parameter problem that is typically present in panel regressions with fixed effects does not affect our results. As shown by Egger, Larch, Staub, and Winkelmann (2011), in a setting with all possible *pairs* of N sectors, the probit model with fixed effects can be estimated consistently. Intuitively, this holds because adding one sector to a dataset with N sectors gives $2N$ additional observations, but only 2 additional fixed effects.

[49]To see this, note that in 1972 (the first year in our panel), $d_{ij}(y-5) = d_{ij}^{67}$. For each benchmark year thereafter, $d_{ij}(y-5)$ reflects the updated input-output network, due to newly formed input connections.

shortest path being semiconductors → electronic components → computing equipment → scales and balances. Thus, when computing equipment adopted semiconductors in 1972, it also reduced the distance between semiconductors and "scales and balances", which eventually adopted semiconductors directly (see also our discussion in Section 2). Our results in Table 4 provide strong evidence that this pattern holds broadly in our data. We find a strong and significant negative relationship between lagged network distance and input adoption in all specifications, and for both the broad and the narrow measure of adoption.[50] Importantly, our results are unchanged when we exclude links within the same 2-digit industry (cols 2 and 5). This makes it unlikely that unobserved trends at the more aggregate industry level are responsible for our results. The coefficient on TFP of input-producing sectors becomes ambiguous and mostly insignificant when we include i-j fixed effects. Thus, network distance turns out to be the more robust among the two main correlates of input adoption in our analysis.

Finally, in Table 5 we use the network distance between i and j following *forward* linkages. We define the forward distance d_{ji} – analogous to the distance based on backward linkages – as the shortest path that connects j and i via *output* flows (beginning from j). For example, if j supplies to k, which in turn supplies to i, then $d_{ji} = d_{jk} + d_{ki}$, where d_{jk} is the shortest-distance forward link between j and k. We find that the coefficient is statistically insignificant and quantitatively small in almost all specifications – this holds irrespective of whether we include only forward distances (cols 1 and 4), or forward distance and backward distance simultaneously. The only time when the coefficient on forward distance is significant (col 3), its sign is positive. Thus, if anything, shorter forward distance is associated with a (marginally) *smaller* probability of adoption. For example, the forward distance from rubber to automobiles is short, with tires as the connecting link. But rubber producers do not adopt cars as an input. These findings make it unlikely that our main measure for network proximity merely captures technologically similar clusters – if this was the case, we should find results irrespective of the direction of input (or output) flows within such clusters. Importantly, our main results are unchanged when we control for forward distance – both for the broad definition (cols 2 and 3) and for the narrow definition (cols 5 and 6) of adoption.

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

[50] The magnitude of the standardized effect is now somewhat smaller than above, at about 0.6% for the broad adoption measure. One reason for this difference may be that we now have to use OLS regressions, which generally yield somewhat lower coefficients than the probit model. We use OLS because, in contrast to our earlier analysis, using probit in a setup with i-j effects would suffer from the incidental parameter problem (see also footnote 48).

sample period in 2002. We run the following regression:

$$T_{ij} = \beta \cdot d_{ij}^{67} + \gamma \cdot \triangle TFP_i + \delta_i + \delta_j + \varepsilon_{ij} \ , \tag{13}$$

where d_{ij}^{67} is network distance in 1967, and $\triangle TFP_i$ denotes the (average annual) change in total factor productivity in the input producing sector between 1967 and the year of adoption.[51] Finally, δ_i and δ_j are input-producing and adopting sector fixed effects, respectively.

Table 6 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. Using also input producing sector fixed effects reflects 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: cols 1-3 do not include δ_i, while all other specifications in Table 6 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 4, with all other variables being the same). In the following, we discuss the individual results in detail.

Column 1 shows a strong positive association between initial network distance and time to adopt. We also find a strong negative relationship between TFP growth in i and average adoption time for i. The main difference with our panel results is that TFP growth now shows a quantitatively important relationship with adoption. A one std increase in $\triangle TFP_i$ is associated with a 1.8 year increase in time to adopt. To put this estimate in context, the average time to adopt (conditional on adoption occurring prior to 2002) in our sample is 16.7 years. One explanation for the larger results on TFP growth is that – in contrast to our panel results in 5-year intervals – the cross-sectional results on time to adopt exploit long-term changes in productivity. Our findings thus suggests that input adoption reacts more to secular trends than to short-time hikes in the efficiency of input production.

In col 2 we use TFP growth of input-producing sectors between to 1958 and 1967, which is highly correlated with the post-1967 TFP growth. Focusing on historical efficiency growth in the input-producing sector i addresses the possibility of reverse causality, i.e., that firms in 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 much smaller than in col 1, with a standardized effect

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

of -0.28 years. To obtain a coefficient estimate that can be more readily compared with our baseline results, and at the same time addresses 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.[52] The first stage has very strong predictive power, with an F-statistic above 800. The 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 -3.88.

For the remaining columns (4-8), we introduce fixed effects also for the input-producing sector. This raises the magnitude of coefficients for both network distance and TFP growth. Our baseline specification in col 4 implies that 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 6.7 years. Columns 5 and 6 show that our results are also robust to excluding early 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 7) and using the narrow definition of adoption (col 8) also yields similar estimates.

6 Firm-Level Evidence

In this section we analyze the relationship between network distance and input adoption at the firm level. This is motivated by the fact that network linkages ultimately reflect the flow of inputs across individual producers. To make progress in this direction, we use some (limited) information on firm-to-firm linkages.

6.1 Description of firm-level data

We use data from Compustat, which includes information on supply linkages. In accordance with Financial Accounting Standards Rule No.131, publicly listed firms are required to disclose the identity of their major customers. A major customer is defined as any firm responsible for more than 10% of the seller's revenues, although firms occasionally report the identity of customers below that threshold. This firm-level network data can be linked to the balance sheet information in Compustat, allowing us to associate information on firms' customers and suppliers with other firm-level observables.[53]

The raw data is reported annually and covers the period from 1977 to 2008 for a total of 43,506

[52]Note that this approach is only feasible if we do not use input producer fixed effects δ_i. Also, it is important to emphasize that we do not interpret this as an 'instrumental variable' regression because it cannot address omitted variable concerns.

[53]Cohen and Frazzini (2008) were the first to explore firm linkages from this data source to examine return predictability across linked firms. Kelly et al. (2013) also use these data to show that firm level volatility depends on the structure of buyer-supplier linkages. Finally, Atalay et al. (2011) use the same data source to develop a model of the buyer-supplier networks in the U.S. economy. We are grateful to the latter set of authors for sharing their data with us.

firm-to-firm links. To reduce noise, we aggregate the information on customer-supplier linkages over non-overlapping 5-year intervals.[54] We define a directed customer-supplier network at the firm level as follows: a directed edge from node i to node j is present if – at any point during the 5-year interval – firm j is reported as a major customer of firm i. Note that this customer-supplier network is binary; weights cannot be computed because information on the value of product flows from i to j is not systematically reported.

Based on this definition of firm-level customer-supplier networks, we define our measures of network distance and input adoption in an analogous way to the sector-level input-output data. First, for each 5-year interval, we define the network distance between any two firms present in the dataset as the length of the shortest directed path between any two nodes. Since the network is binary, distance reflects the minimum *number* of directed edges that lead from i to j. If i supplies directly to j, distance $d_{ij} = 1$; if i supplies to k, and k to j, then $d_{ij} = 2$, etc. Second, we say that firm j has adopted an input being supplied by firm i if, in a 5-year interval, firm i reports firm j as a major customer, and it did not do so at any previous time in our dataset. Our final panel dataset includes approximately 14.5 million firm pairs with distance $d_{ij} > 1$, i.e., firms that are not directly linked. About 1,200 firm pairs have distance $d_{ij} = 2$; 200 have $d_{ij} = 3$, and a few have distances 4 or 5. For the vast majority of firm pairs, no path exists based on the binary network (so that $d_{ij} = \infty$). The lack of network connections is in part due to the restrictive nature of the data with the 10% reporting threshold. We thus interpret our results as exploratory rather than conclusive.

As mentioned above, we supplement these data with other firm-level observables available from Compustat. These include firm employees and sales. As a proxy for firm-level productivity growth, we first compute labor productivity as sales per worker and derive its growth rate over each 5-year window. In order to control for geographical proximity between two firms, we compute the distance between their headquarters. To proxy for technological similarity, we use the 4-digit SIC code classifying the main sector of activity of each firm.

6.2 Firm-level results

For all i-j pairs that were not directly connected in any 5-year period prior to y, we run panel regressions of the form:

$$Prob\left(A_{ij}(y) = 1\right) = g\left(I_{ij}(y-5), X_i(y), X_j(y)\right) \ , \tag{14}$$

where $A_{ij}(y)$ is an indicator that equals one if firm j adopted input i in the 5-year period y. Our main explanatory variable is an indicator for whether firms i and j were indirectly connected, via

[54] The 5-year intervals are 1977-81, 1982-86,...,2002-06.

one other node, in the previous period: the dummy $I_{ij}(y-5)$ equals one if the binary directed distance between i and j equalled 2. The coefficient on this variable reflects by how much the probability of adoption in any given interval increases if $d_{ij} = 2$, as compared to $d_{ij} > 2$.[55] X_i (X_j) are additional controls for the input-producing (adopting) firm, such as the number of employees, output per worker, and the (lagged) growth thereof. We also include fixed effects for each time period, and for adopting and producing firms. Due to the large number of firm pairs, the probit and the hazard model are computationally unfeasible in the presence of firm fixed effects. We thus use the linear probability model for the functional form of $g(\cdot)$ in those specifications.

Table 7 presents our firm-level results. In column 1 we use a simple OLS regression with time-period dummies, including only $I_{ij}(y-5)$ and the geographic distance between firms. Both are highly significant and have the expected sign – previous indirect connections (i.e., network proximity) increase the probability of new link formation, while geographic distance reduces this probability. To interpret and compare the magnitude of our estimates, we provide standardized coefficients in square brackets. For our main explanatory variable, these reflect the change in adoption probability when an indirect link existed ($I_{ij}(y-5) = 1$), as compared to when it did not exist ($I_{ij}(y-5) = 0$). For all continuous variables, the standardized coefficients represent the change in adoption probability due to a one standard deviation increase in the explanatory variable. The standardized coefficient for network proximity is economically sizeable, with 2.85%. In contrast, geographical distance has a minuscule effect: reducing log distance by one standard deviation increases the probability of j adopting i by merely 0.007%. Both the magnitude and significance of these effects are confirmed by the probit model in column 2.

In column 3 we introduce fixed effects for input-producing and (potential) adopting firms, i and j. Our results remain unchanged. The same is true when we add controls for firm size, productivity, and productivity growth (column 4). The coefficient for productivity growth in input-producing firms is positive and significant. This confirms the sector-level result that efficiency gains in input production are associated with a higher likelihood of adoption (however, the effect is quantitatively small). In column 5 we exclude all $i-j$ firm pairs that belong to the same 2-digit SIC industry. The coefficient on $I_{ij}(y-5)$ falls only slightly and remains highly significant. This implies that technological proximity alone is probably not the main driver of our results – initial network proximity raises the likelihood of adoption also for inputs from different industries. Finally, we restrict the sample to input-producing firms i from manufacturing (col 6), and from the service sector (col 7). We find very similar results in both samples, suggesting that the role of network proximity is not limited to physical inputs.

[55] As explained above, the vast majority of the remaining firm pairs have an infinite binary distance. All results are practically identical when we exclude the few cases with $d_{ij} =$ 3,4, or 5.

7 Conclusion and Broader Implications

Input-output linkages have important implications for macroeconomic outcomes. While typically observed at the sectoral level, these linkages reflect the flow of products between individual producers, and thus ultimately the underlying technology at the product level. The evolution of the input-output structure is therefore at the heart of technological progress. We studied the mechanism of input link formation both theoretically and empirically. Guided by a stylized model of directed search in a network, we uncovered a strong and novel empirical regularity: sectors (and firms) that are closer in the input-output network are significantly more likely to form new input linkages. In other words, the existing production network plays an important role for the diffusion of inputs and thus for the evolution of the input-output network itself.

Our theoretical and empirical results have several important implications. First, from a network perspective, General Purpose Technologies (GPTs) correspond to central nodes, i.e., extremely prominent inputs. Our findings shed new light on the rise of general purpose inputs: inputs that are initially closer to many potential adopters are more likely to become widely adopted. The left panel of Figure 6 illustrates this finding. It plots the (log) number of sectors that adopted input i after 1967 against the initial average network distance of i (see note to figure for details). The latter is low whenever i is *indirectly* linked to many, relatively large, manufacturing sectors; it thus reflects the "network growth potential" of input i. In the regression underlying the figure, network growth potential accounts for more than 25% of cross-sectional variation in input diffusion (based on the R^2).[56] In sum, the figure shows that network proximity is a crucial determinant of input diffusion, and is thus a potentially important factor in the rise of GPTs.

Second, our findings can help to explain sector-level growth patterns. Intuitively, if initial network proximity is associated with extensive margin growth (i.e., input diffusion), variation in the former should also predict differential growth across sectors. The right panel of Figure 6 shows that this is indeed the case. It plots 1967-2002 employment growth for each input-producing sector against our measure of initial "network growth potential". This relationship is naturally more noisy than the result on input diffusion, because growth is affected by many drivers other than adoption by other sectors. Nevertheless, the result is quantitatively important: a one-standard deviation increase in initial network distance is associated with a (highly significant) decline in sector-level growth by 0.2 standard deviations. This suggests that our network view of input diffusion may have implications for structural change and unbalanced growth (see Herrendorf, Rogerson, and Valentinyi, 2014, and references therein).

[56] We note in passing that this dwarfs the explanatory power of an input producer's TFP growth, which accounts for merely 1% (despite the fact that it is statistically highly significant).

Third, our results give rise to a possible new channel by which misallocation may affect aggregate productivity. Input-output linkages can amplify micro-level distortions to a *given* sector i, leading to *static* losses in aggregate efficiency (Jones, 2013). In our setting, distortions can also give rise to *dynamic* aggregate productivity losses: distorting sectors that use i as an input can affect i's subsequent diffusion. To see this, consider our earlier example of semiconductors. A crucial gateway that connected these to other sectors in the economy was the "Electronic Components" sector. Consequently, (hypothetical) distortions to the latter could have stunted the diffusion of semiconductors. Thus, our findings suggest the possibility that distortions to "network bottlenecks" can have an impact on aggregate productivity growth.

These broader implications underline the need to shed light on the mechanisms behind our findings. We have discussed technological proximity, coagglomeration, and information diffusion along input linkages as possible explanations for the strong relationship between network proximity and input diffusion. Our findings suggest that *vertical* distance along supply chains is a promising starting point to understand patterns of input adoption and diffusion. In this paper, we have taken the necessary first step of documenting this novel pattern in the data; we leave the systematic assessment of the underlying mechanisms for future research.

References

Acemoglu, D., U. Akcigit, and W. Kerr (2014). The Innovation Network. Working Paper.

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.

Antràs, P. and D. Chor (2013). Organizing the Global Value Chain. *Econometrica 81*(6), 2127–2204.

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

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. Nirei, Y. Saito, and A. Tahbaz-Salehi (2014). Supply Chain Disruptions: Evidence from the Great East Japan Earthquake. mimeo.

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

Carvalho, V. M. (2014). From Micro to Macro via Production Networks. *Journal of Economic Perspectives 28*(4), 23–48.

Chaney, T. (2014). The Network Structure of International Trade. *American Economic Review 104*(11), 3600–34.

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

Cohen, L. and A. Frazzini (2008). Economic Links and Predictable Returns. *Journal of Finance 63*(4), 1977–2011.

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.

Costinot, A., J. Vogel, and S. Wang (2013). An Elementary Theory of Global Supply Chains. *Review of Economic Studies 80*(1), 109–144.

di Giovanni, J., A. A. Levchenko, and I. Mejean (2014). Firms, Destinations, and Aggregate Fluctuations. *Econometrica 82*(4), 1303–1340.

Dorogovtsev, S., J. Mendes, and A. Samukhin (2000). Structure of Growing Networks with Preferential Linking. *Physical Review Letters 85*, 4633–6.

Egger, P., M. Larch, K. E. Staub, and R. Winkelmann (2011). The Trade Effects of Endogenous Preferential Trade Agreements. *American Economic Journal: Economic Policy 3*(3), 113–43.

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.

Fafchamps, M., S. Goyal, and M. J. van der Leij (2010). Matching and Network Effects. *Journal of the European Economic Association 8*(1), 203–231.

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.

Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review 98*(1), 394–425.

Garcia-Marin, A. and N. Voigtländer (2013). Exporting and Plant-Level Efficiency Gains: It's in the Measure. NBER working paper #19033.

Ghiglino, C. (2012). Random Walk to Innovation: Why Productivity Follows a Power Law. *Journal of Economic Theory 147*(2), 713–737.

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

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

Herrendorf, B., R. Rogerson, and Ákos. Valentinyi (2014). *Growth and Structural Transformation*, Volume 2 of *Handbook of Economic Growth*, Chapter 6, pp. 855–941. Elsevier.

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. (2011). Intermediate Goods and Weak Links in the Theory of Economic Development. *American Economic Journal: Macroeconomics 3*(2), 1–28.

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.

Oberfield, E. (2013). Business Networks, Production Chains and Productivity: A Theory of Input-Output Architecture. Working Paper, Princeton University.

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.

Weitzman, M. L. (1998). Recombinant Growth. *Quarterly Journal of Economics 113*(2), 331–360.

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.

FIGURES

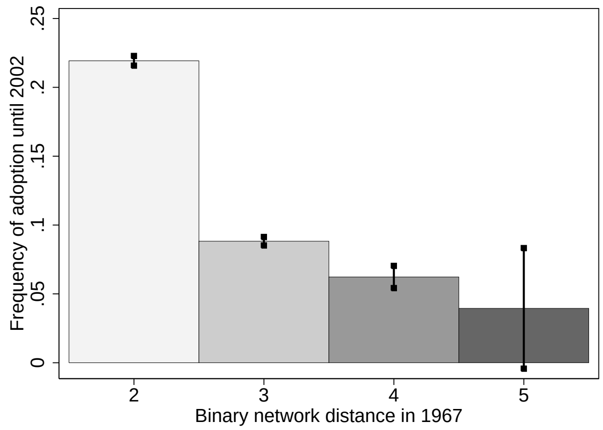

Figure 1: Input adoption and initial network distance

Notes: The figure shows that sector pairs that are closer in the U.S. input-output network in 1967 are more likely to see direct adoption by 2002. The x-axis shows the binary distance between input i and a potential adopting sector j in 1967. Sectors that are already directly connected in 1967 (distance 1) are excluded from the analysis. The y-axis shows the share of all sector pairs for which j as adopted i by 2002. The whiskers correspond to the 95% confidence intervals.

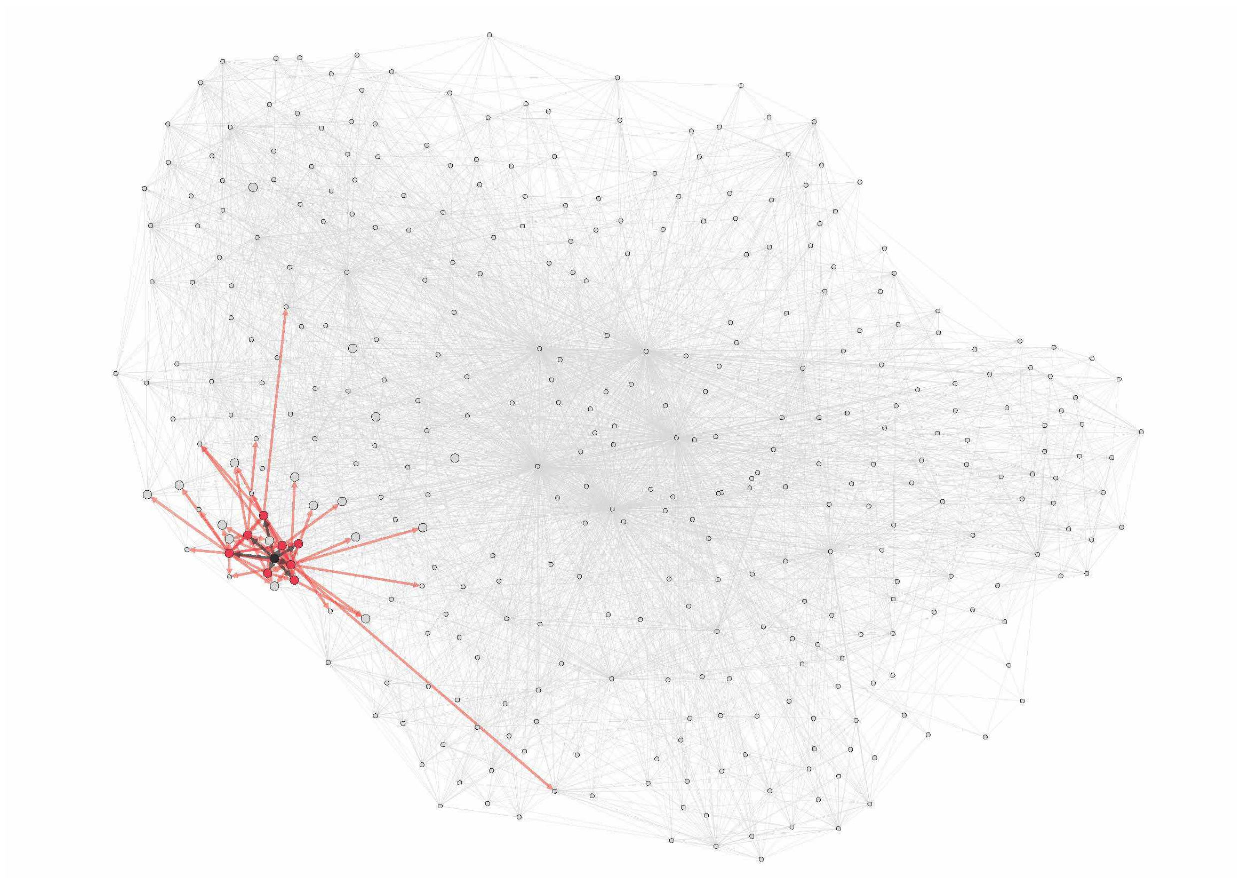

Figure 2: Input-output network and semi-conductor linkages in 1967

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.

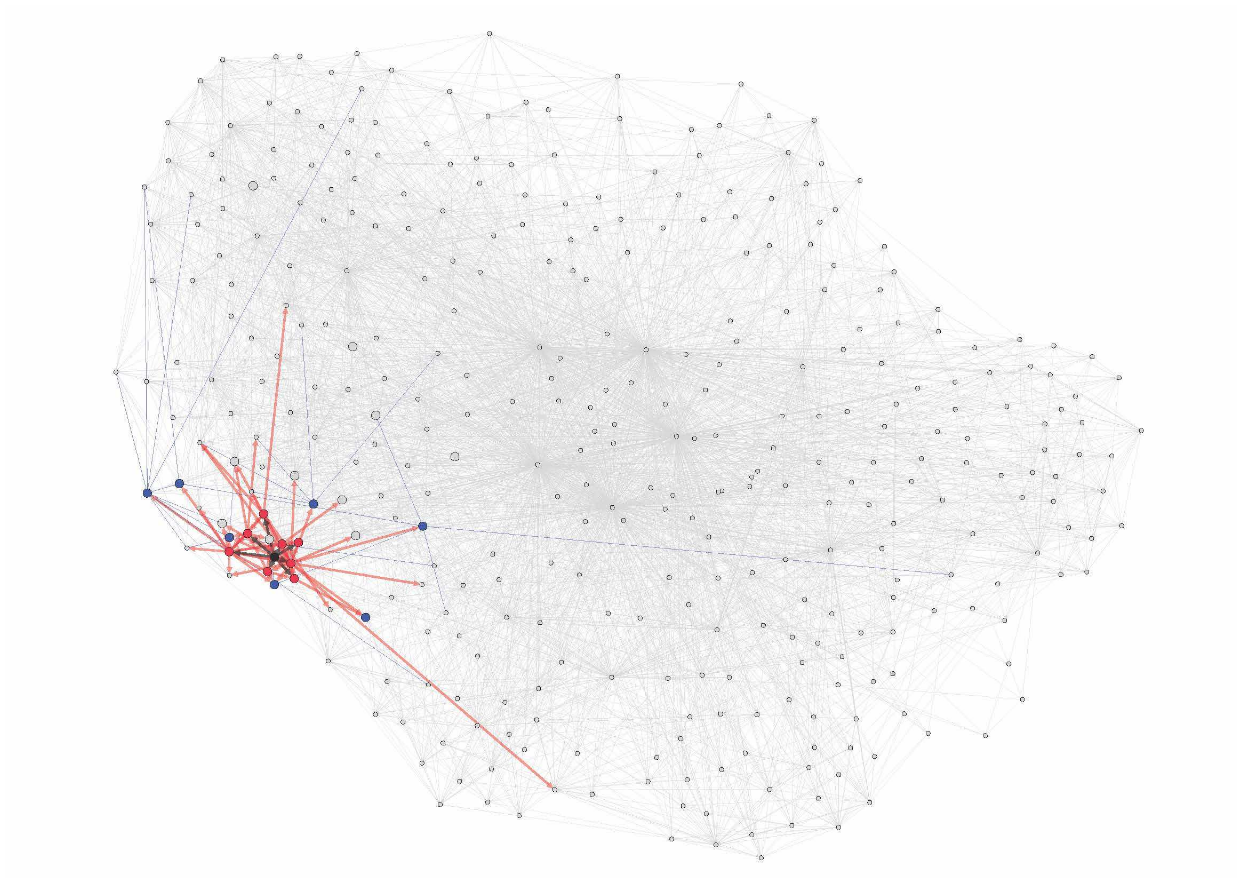

Figure 3: Adoption of semi-conductors in 1972

Notes: The figure shows the U.S. input-output network in 1972. The dots and arrows are the same as in Figure 2. In addition, blue dots represent sectors adopting semiconductors in 1972.

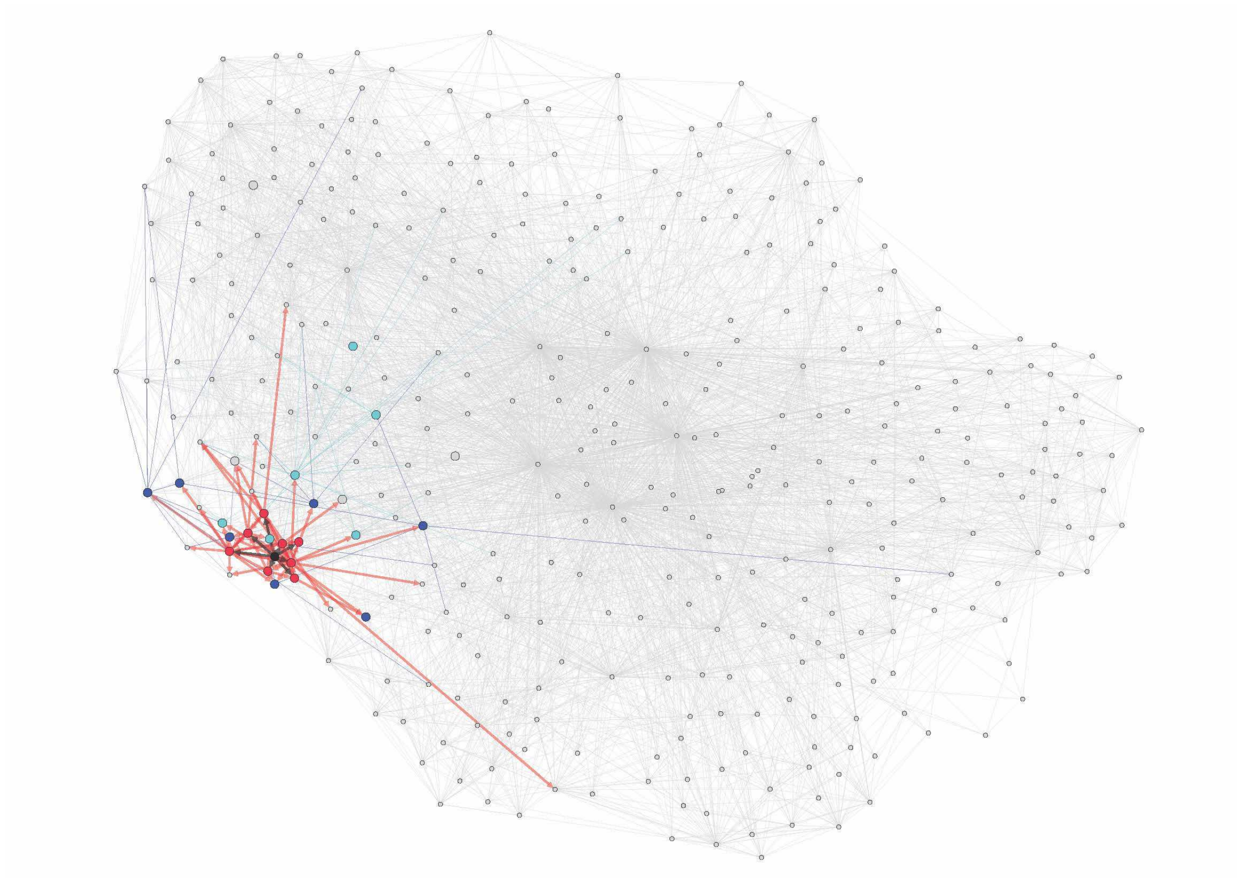

Figure 4: Adoption of semi-conductors in 1977

Notes: The figure shows the U.S. input-output network in 1977. The dots and arrows are the same as in Figure 2. In addition, blue dots represent sectors adopting semiconductors in 1972, and cyan dots are sectors adopting semiconductors in 1977.

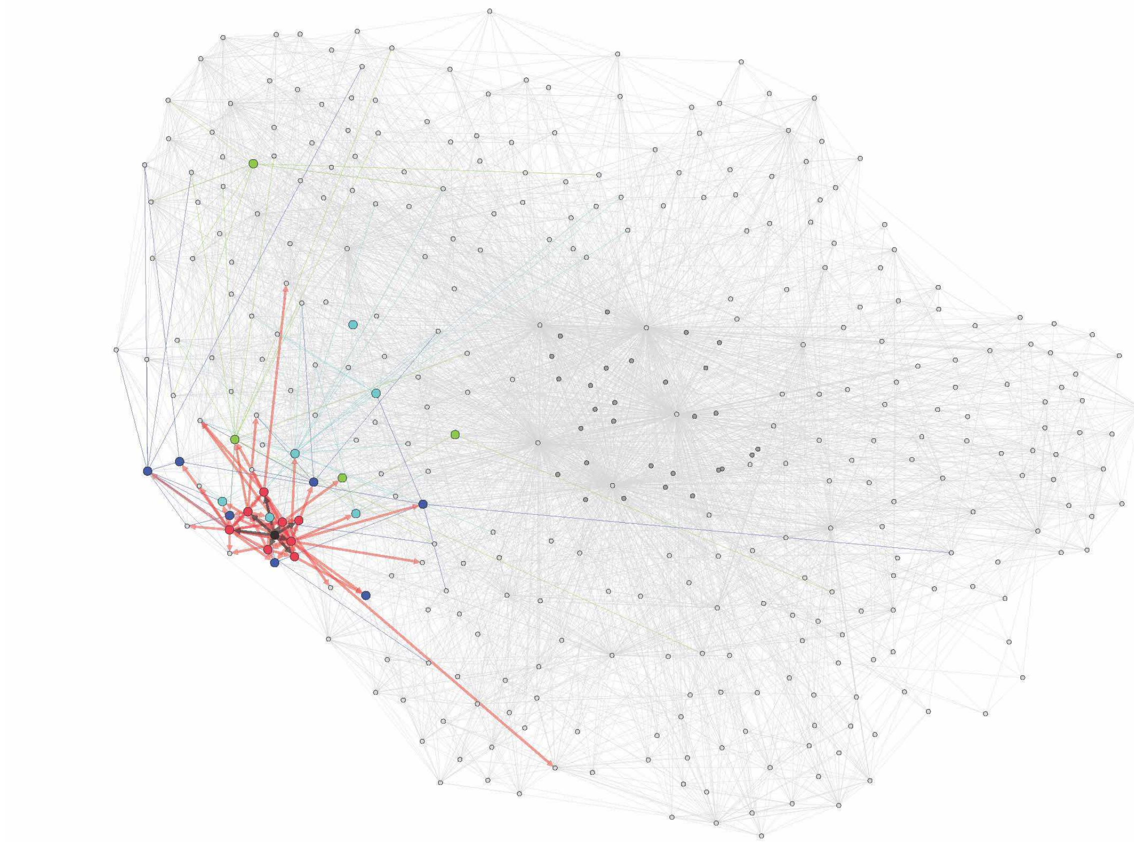

Figure 5: Adoption of semi-conductors in 1982

Notes: The figure shows the U.S. input-output network in 1982. The dots and arrows are the same as in Figure 2. 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 6: Initial network distance, input diffusion, and employment growth

Notes: The left panel shows the relationship between an input sector's (i) initial average network distance in 1967 and its subsequent adoption by other sectors (j) over the period 1967 and 2002. The right panel shows the relationship between initial average distance and subsequent employment growth of input sector i. The x-axes of both panels display the average network distance of an input i to all other sectors (j) in 1967. This measure is computed as follows: first, we calculate our baseline network distance measure (weighted and directed) for each $i - j$ pair that is not yet directly connected in 1967. For each input i, we then compute the weighted average over all sectors j, where weights are given by the total value of sector j's output, relative to aggregate manufacturing output in 1967. The y-axis of the left panel gives the (log of) total number of sectors j which adopted i as an input in subsequent years. We use our most conservative notion of adoption, by requiring that i supplies to j no less than a \$1million for at least 15 years after the initial adoption date. The y-axis in the right panel gives the 1967-2002 employment growth for each sector. Regressions in both panels control for (log) initial employment in sector i and for its TFP growth between 1967 and 2002.

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 $d_{ij}(y-5)$	-0.1906***	-0.1904***	-0.0089***	-0.0089***	0.5940***	0.5946***
	(0.0039)	(0.0039)	(0.0002)	(0.0002)	(0.0055)	(0.0055)
	[-2.31%]	*[-2.31%]*	*[-1.41%]*	*[-1.41%]*	*[-4.12%]*	*[-4.18%]*
$\triangle_5 TFP_i$		0.0623*		0.0119***		1.405***
		(0.0372)		(0.0024)		(0.1038)
		[0.04%]		*[0.09%]*		*[0.13%]*
$\triangle_5 TFP_j$		-0.0376		-0.0020		1.3463
		(0.1420)		(0.0107)		(0.4095)
Observations	563,173	563,173	563,173	563,173	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. 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.

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)	(7)	(8)
Links excluded‡		2-digit				2-digit		
	Broad definition of adoption				Narrow definition of adoption			
Distance $d_{ij}(y-5)$	-0.162***	-0.158***		-0.104***	-0.280***	-0.278***		-0.134***
	(0.005)	(0.005)		(0.005)	(0.008)	(0.009)		(0.009)
	[-1.83%]	*[-1.73%]*		*[-1.07%]*	*[-1.74%]*	*[-1.68%]*		*[-0.75%]*
Distance in 1967, d_{ij}^{67}			-0.219***	-0.100***			-0.397***	-0.253***
			(0.010)	(0.011)			(0.012)	(0.015)
			[-0.94%]	*[-0.44%]*			*[-0.93%]*	*[-0.31%]*
TFP_i	-0.192***	-0.152***	-0.091**	-0.113**	1.072***	1.196***	1.287***	1.301***
	(0.040)	(0.043)	(0.043)	(0.044)	(0.068)	(0.065)	(0.066)	(0.072)
	[0.27%]	*[0.27%]*	*[0.28%]*	*[-0.10%]*	*[0.08%]*	*[0.06%]*	*[0.02%]*	*[-0.13%]*
TFP_j	-0.019	-0.036	0.037	0.036	0.061	0.075	0.080	0.081
	(0.115)	(0.116)	(0.130)	(0.129)	(0.097)	(0.102)	(0.108)	(0.109)
$\triangle_5 TFP_i$	0.251***	0.261***	0.280***	0.276***	0.142***	0.122***	0.046*	0.034
	(0.015)	(0.015)	(0.015)	(0.015)	(0.024)	(0.024)	(0.024)	(0.025)
$\triangle_5 TFP_j$	-0.031	-0.021	-0.009	-0.010	0.013	0.018	0.032	0.033
	(0.048)	(0.051)	(0.052)	(0.052)	(0.041)	(0.044)	(0.042)	(0.042)
$\ln(emp)_i$	0.126***	0.137***	0.101***	0.102***	0.176***	0.192***	0.134***	0.135***
	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)
$\ln(emp)_j$	0.089***	0.094***	0.096***	0.096***	0.159***	0.164***	0.174***	0.174***
	(0.007)	(0.007)	(0.008)	(0.008)	(0.010)	(0.010)	(0.011)	(0.011)
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	519,041	611,669	430,836	307,278	380,573	380,573	309,499	291,397

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	1997	1997	1997	none	1997	1997
Links excluded‡			2-digit				2-digit	
	Broad definition of adoption				Narrow definition of adoption			
PANEL A: All input relationships								
Distance $d_{ij}(y-5)$	-0.208***	-0.145***	-0.199***	-0.062***	-0.362***	-0.362***	-0.346***	-0.163***
	(0.012)	(0.008)	(0.012)	(0.014)	(0.024)	(0.024)	(0.026)	(0.021)
	[-1.52%]	*[-1.51%]*	*[-1.52%]*	*[-0.36%]*	*[-1.19%]*	*[-1.19%]*	*[-1.25%]*	*[-0.47%]*
Distance in 1967, d_{ij}^{67}				-0.600***				-0.669***
				(0.035)				(0.039)
				[-1.64%]				*[-1.07%]*
TFP_i	0.165***	0.195***	0.190***	-0.075	0.085	0.085	0.037	0.126*
	(0.034)	(0.029)	(0.037)	(0.060)	(0.065)	(0.065)	(0.070)	(0.073)
	[0.14%]	*[0.21%]*	*[0.16%]*	*[-0.04%]*	*[0.04%]*	*[0.04%]*	*[0.02%]*	*[0.05%]*
Control Variables§	✓	✓	✓	✓	✓	✓	✓	✓
Using Sector FE	✓	✓	✓	✓	✓	✓	✓	✓
Producing Sector FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	519,041	611,669	430,836	307,278	380,573	380,573	309,499	291,397
PANEL B: Exclude links that reflect less than $1million input purchase								
Distance $d_{ij}(y-5)$	-0.173***	-0.138***	-0.161***	-0.063***	-0.231***	-0.231***	-0.210***	0.022
	(0.014)	(0.012)	(0.014)	(0.014)	(0.027)	(0.027)	(0.029)	(0.022)
	[-0.92%]	*[-0.88%]*	*[-0.86%]*	*[-0.30%]*	*[-0.62%]*	*[-0.62%]*	*[-0.57%]*	*[0.05%]*
Distance in 1967, d_{ij}^{67}				-0.629***				-0.829***
				(0.031)				(0.038)
				[-1.67%]				*[-1.36%]*
TFP_i	0.084**	0.199***	0.088*	0.062	0.089	0.089	0.047	0.055
	(0.041)	(0.038)	(0.046)	(0.058)	(0.062)	(0.062)	(0.068)	(0.067)
	[0.05%]	*[0.14%]*	*[0.05%]*	*[0.03%]*	*[0.03%]*	*[0.03%]*	*[0.02%]*	*[0.02%]*
Controls as in Panel A	✓	✓	✓	✓	✓	✓	✓	✓
Observations	577,122	671,697	482,850	375,841	398,169	398,169	323,867	324,315

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. For description explanatory variables and additional detail see the note to Table 2. 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 4 and 8 exclude all i-j pairs that belong to the same 2-digit industry.

§ Control Variables include TFP in the input-adopting industry, and (log) employment in both adopting and input-producing industries.

Table 4: Panel regressions with pairwise fixed effects

Dep. Var.: Dummy for adoption of input i by sector j in year y	(1)	(2)	(3)	(4)	(5)	(6)
Remarks		2-digit[†]	1million[‡]		2-digit	1million[‡]
	Broad definition of adoption			Narrow definition of adoption		
Distance $d_{ij}(y-5)$	-0.0040***	-0.0037***	-0.0041***	-0.0018***	-0.0017***	-0.0016***
	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0001)	(0.0001)
	[-0.64%]	*[-0.59%]*	*[-0.57%]*	*[-0.28%]*	*[-0.28%]*	*[-0.22%]*
TFP_i	0.0034*	0.0029	-0.0107***	0.0002	-0.0003	-0.0040***
	(0.0019)	(0.0019)	(0.0015)	(0.0008)	(0.0007)	(0.0006)
	[0.05%]	*[0.04%]*	*[-0.16%]*	*[0.00%]*	*[-0.00%]*	*[-0.06%]*
Controls[§]	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Pairwise i-j FE	✓	✓	✓	✓	✓	✓
Observations	556,936	523,627	568,771	555,529	522,718	566,600

Notes: All regressions are estimated by OLS. 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. TFP_i is the level of TFP in the input-producing industry in year y. 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.

‡ Columns 3 and 6 exclude links that reflect less than $1million input purchase.

§ Controls include TFP in the input-adopting industry, and employment in both adopting and input-producing industries.

Table 5: Panel results: Forward linkage distance

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		
Forward Distance $d_{ji}(y-5)$	0.010	0.010	0.027**	-0.014	-0.015	0.015
	(0.012)	(0.012)	(0.012)	(0.013)	(0.013)	(0.012)
	[0.09%]	*[0.08%]*	*[0.25%]*	*[-0.08%]*	*[-0.07%]*	*[0.08]*
Distance $d_{ij}(y-5)$		-0.213***	-0.206***		-0.361***	-0.350***
		(0.012)	(0.013)		(0.024)	(0.026)
		[-1.56%]	*[-1.56%]*		*[-1.22%]*	*[-1.29%]*
Controls	✓	✓	✓	✓	✓	✓
Observations	488,947	488,947	406,865	358,171	358,171	292,034

Notes: 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. All regressions are estimated by Probit. Controls are all those used in Table 3, Panel A (including all fixed effects). "Forward distance $d_{ji}(y-5)$" is network distance (with a 5-year lag), using forward linkages from sector j to sector i, i.e., via other sectors that j supplies to. For description of the remaining explanatory variables and additional detail see the note to Table 2. 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 3 and 6 exclude all i-j pairs that belong to the same 2-digit industry.

Table 6: Time to adoption

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Years excluded	1997	1997	1997	1997	1972,97	none	1997	1997
Other remarks			2SLS§				2-digit†	narrow‡
Distance d_{ij} in 1967	0.937***	0.968***	0.976***	3.112***	1.778***	3.104***	3.307***	1.228***
	(0.196)	(0.212)	(0.182)	(0.341)	(0.360)	(0.311)	(0.354)	(0.290)
	[0.64]	*[0.66]*	*[0.67]*	*[2.14]*	*[1.15]*	*[2.04]*	*[2.28]*	*[0.72]*
$\triangle TFP_i(1967 - y_{adopt})$	-96.925***		-211.401***	-364.787***	-331.477***	-281.502***	-376.759***	-146.929***
	(3.919)		(24.137)	(13.186)	(26.434)	(11.861)	(14.029)	(12.575)
	[-1.78]		*[-3.88]*	*[-6.70]*	*[-3.95]*	*[-4.37]*	*[-6.97]*	*[-3.06]*
$\triangle TFP_i(1958 - 67)$		-18.341***						
		(6.189)						
		[-0.28]						
Using Sector FE	✓	✓	✓	✓	✓	✓	✓	✓
Producing Sector FE				✓	✓	✓	✓	✓
R^2	0.19	0.17	0.16	0.73	0.72	0.67	0.73	0.66
Observations	14,849	15,072	14,849	14,849	8,604	24,312	13,856	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 (11). 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.

§ 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.

† 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 7: Firm level panel results

Dep. Var.: Dummy for firm j adopting inputs from firm i in a 5-year time interval y

	(1) OLS	(2) Probit	(3) OLS	(4) OLS	(5) OLS	(6) OLS	(7) OLS
Sample					2-digit†	Manufacturing	Services
$I_{ij}(y-5)$	0.02854***	1.61359***	0.02161***	0.02140**	0.01834**	0.01966**	0.02367*
	(0.00745)	(0.11780)	(0.00779)	(0.00888)	(0.00806)	(0.00904)	(0.01348)
	[2.85%]	*[2.69%]*	*[2.16%]*	*[2.14%]*	*[1.83%]*	*[1.97%]*	*[2.37%]*
ln(geodistance)	-0.00006***	-0.05809***	-0.00007***	-0.00007***	-0.00006***	-0.00006***	-0.00007***
	(0.00001)	(0.00604)	(0.00001)	(0.00001)	(0.00001)	(0.00002)	(0.00001)
	[-0.007%]	*[-0.005%]*	*[-0.007%]*	*[-0.007%]*	*[-0.006%]*	*[-0.006%]*	*[-0.007%]*
$\triangle_5 \ln(Y/L)_i$				0.00003***	0.00003***	0.00003**	0.00003
				(0.00001)	(0.00001)	(0.00001)	(0.00002)
Controls				✓	✓	✓	✓
Using Firm FE			✓	✓	✓	✓	✓
Producing Firm FE			✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓
Observations	14,634,939	14,634,939	14,634,939	8,895,481	8,461,685	4,906,536	3,381,959

Notes: The dependent variable is a dummy that takes on value 1 if firm j adopts input i in a given 5-year interval y between 1977 and 2006. $I_{ij}(y-5)$ is an indicator that equals one if firms i and j were indirectly linked (had a binary distance of 2) in the previous five-year interval. The variable geodistance is the geographical distance between i and j. $\triangle_5 \ln(Y/L)_i$ denotes the change in output per worker in the input-producing firm (i) over the previous (lagged) 5-year interval. Controls include the change in output per worker in the input-using firm over the previous 5 year interval ($\triangle_5 \ln(Y/L)_j$), as well as output per worker and ln(employment) for *both* input-producing and input-using firms. For a description of the firm-level dataset see Section 6.1. Standard errors in parentheses, clustered at the adopting firm (j) level. * p<0.1, ** p<0.05, *** p<0.01. Values in *[square brackets]* are standardized coefficients. For the dummy $I_{ij}(y-5)$ (all other explanatory variables), these reflect the change in adoption probability due to an increase from 0 to 1 (a one standard deviation increase in the explanatory variable).

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

Online Appendix

Input Diffusion and the Evolution of Production Networks

Vasco M. Carvalho
Cambridge and CEPR

Nico Voigtländer
UCLA, NBER, and CEPR

A Proofs and Additional Detail on the Model

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.[1] 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:

$$p_N \left(k_{s_j} \frac{i_{s_j}}{xm} + (m_K - k_{s_j}) \frac{d_i^{out}(t)}{t} \right) \frac{m_N}{m_K m} \tag{A.1}$$

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 (A.1) 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.

[1] 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.

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

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

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:

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

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 $\sum_{i' \in s_i} i'_{s_j} = \mu'_{s_j} \nu_{s_i} \equiv n_{(s_i, s_j)}$. Thus, if s_i is closer to s_j than to $s_{j'}$ at time $t-1$, then s_j is more likely to adopt from s_i at time t, as claimed in the proposition. $\square$

Proof of Proposition 3

Proof. 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, and it is given by $1/J$. 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. Therefore, the number of varieties classified into sector s_i at time t, $K_{s_i}(t)$, is given by a Binomial distribution, $B(t, \frac{1}{J})$.

Second, under the assumption of price symmetry, the sectoral (weighted) outdegree, $w^{out}_{s_i}(t)$ at time t, is proportional to the total number of varieties to which a sector s_j supplies inputs at time t, where the constant of proportionality is given by the price ϕ.[2] Thus

$$w^{out}_{s_i}(t) = \sum_{k=1}^{K_{s_i}(t)} \phi d^{out}_i(t)$$

where $K_{s_i}(t)$ is the (random) number of varieties classified into sector s_i at time t.

Third, we are given that the variety-level outdegree, $d^{out}_i(t)$, is power law distributed. Notice further that, since $K_{s_i}(t)$ is distributed as a Binomial distribution, we have that:

$$\text{Prob}(K_{s_i}(t) > x) = o(\text{Prob}(d^{out}_i(t) > x)$$

that is, $\lim_{x\to\infty} \frac{Prob(K_{s_i}(t)>x)}{Prob(d^{out}_i>x)} = 0$. This is immediate from the fact that power law distributions are heavy-tailed while binomials are thin tailed.

Given the above observations, we can apply known results regarding the tail behavior of random sums of power-law distributed variables. From Lemma 3.7.(1) in Jessen and Mikosch (2006, p.8) we conclude that, as $x \to \infty$

$$\text{Prob}(w^{out}_{s_i}(t) > x) \sim \phi E_t(K_{s_i}(t))\text{Prob}(d^{out}_i(t) > x) = \frac{\phi t}{J}\text{Prob}(d^{out}_i(t) > x)$$

That is, we have shown that $w^{out}_{s_i}(t)$, the sectoral weighted outdegree of sector s_i, inherits the outdegree tail behavior of the varieties classified into it. Since all sectors are simply random collections of varieties with the same outdegree distribution, this result holds true for every sector. Therefore, we have shown that if the variety-level outdegree distribution at time t is power law, so is the distribution of sectoral weighted outdegrees. $\square$

[2] The weighted outdegree refers to values of input flows, while our variety-level predictions 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.

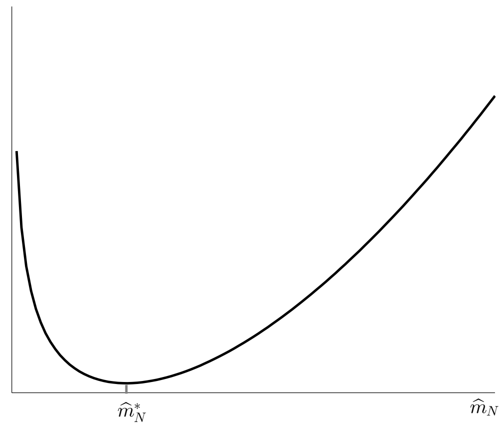

Figure A.1: Optimal choice of network input adoption

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 $\widehat{m}_N^*$.

References

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

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