Bottlenecks: Sectoral Imbalances and the US Productivity Slowdown*

Daron Acemoglu
MIT and NBER

David Autor
MIT and NBER

Christina Patterson
Chicago Booth and NBER

January 2022

Abstract

Despite the rapid pace of innovation in information and communications technologies (ICT) and electronics, aggregate US productivity growth has been disappointing since the 1970s. We propose and empirically explore the hypothesis that this is because of the unbalanced sectoral distribution of innovation over the last several decades. Because an industry's success in innovation depends on complementary innovations among its input suppliers, rapid productivity growth in just a subset of sectors may create bottlenecks and fail to translate into commensurate aggregate productivity gains. Using data on input-output linkages, citation linkages, industry productivity growth and patenting, we find evidence in support of this hypothesis: the variance of supplier TFP growth or innovation adversely affects an industry's own TFP growth and innovation. Our estimates suggest that a substantial share of the productivity slowdown in the US and several other industrialized economies can be accounted for by the sizable increase in cross-industry variance of TFP growth and innovation. For example, if the TFP growth variance had remained at the same level as between 1977 and 1987, the US manufacturing productivity would have grown twice as fast in 1997-2007 as it did, reaching a higher level of growth than in 1977-1987 and 1987-1997.

Keywords: innovation, input-output linkages, productivity.
JEL Classification: O47, O30.

*We thank Nick Bloom, Steve Davis, Bill Kerr and John Van Reenen for helpful comments and Rebecca Jackson and Zhe Fredric Kong for expert research assistance. Acemoglu and Autor acknowledge support from Accenture LLP, IBM Global Universities, Schmidt Futures, and the Smith Richardson Foundation. Acemoglu also acknowledges support from Google, the National Science Foundation, the Sloan Foundation, and the Toulouse Network on Information Technology. Autor also thanks the Carnegie Fellows Program, the Heinz Endowments, and the Washington Center for Equitable Growth.

1 Introduction

One of the most enduring macroeconomic puzzles of the last several decades is the pervasive slowdown in productivity growth across industrialized nations in the midst of breakneck advances in information and communication technologies (ICT) and electronics. Figure 1 provides a glimpse of recent breakthroughs in ICT and electronics by plotting the distribution of patents granted over the last several decades.[1] Two patterns are evident from the figure: first, there is a rapid takeoff in the number of total patents in the 1980s, and second, the share of ICT and electronics patents surges during the same time interval. For example, between 1990 and 2010, the total number of patents granted annually rose from 99,000 to 208,000, while the number of ICT and electronics patents granted increased by approximately 87,000, accounting for the bulk of increased patenting in the US. Figure 2 depicts the growth rate of Total Factor Productivity (TFP) in the US economy and in the leading OECD economies in recent decades. We see that productivity growth has been slower everywhere, with the possible exception of Germany. In the US, for example, there has been very little TFP growth since 2005.

The exponential advance of ICT and electronics innovation has led some commentators to conclude that we are on the verge of a new age of abundance, or even "singularity" driven by "superintelligent" machines (e.g., Kurzweil (2005), Diamandis and Kotler (2012), and Bostrom (2014)). Others, looking at the TFP data, have concluded that we have entered an age of slower growth because most promising technologies have already been developed and exploited (Cowen (2011), Gordon (2017)).[2]

This paper offers a potential reconciliation of these two puzzling macroeconomic trends based on the idea that technological advances over the last several decades have been unbalanced across sectors and have thus created endogenous bottlenecks, holding aggregate productivity back. We propose a simple framework in which new technologies or products in a sector require simultaneous improvements in the quality of the several inputs. For example, breakthroughs in automotive technology cannot be achieved just with advances in sensors and

[1] The orange bars correspond to the share of USPTO patents granted in ICT and electrical products and electronics, while the blue bars plot the share patents granted in electricity and electronics. The green line shows the total number of patents granted. See below for data details.

[2] Those subscribing to the first view often claim that growth is mismeasured, which is undoubtedly true. Nevertheless, mismeasurement does not seem to account for the broad outlines of the productivity slowdown since the 1970s. First, growth was almost surely also mismeasured in the decades that followed World War II, when many new consumer goods and technologies were introduced. Second, many of the implications of growth mismeasurement thesis, such as faster productivity growth in sectors with less mismeasurement potential, do not receive support from the data (Byrne et al. (2016), Syverson (2017)). Third, there is no evidence for even the most basic predictions of the ICT-driven fast productivity growth, for example, with more ICT-intensive industries (outside of ICT-producing industries themselves) exhibiting, if anything, *slower* growth of *nominal* and *real* value-added (Acemoglu et al. (2014)).

software, but will necessitate compatible improvements in energy storage, drivetrains, and tire adhesion. When some of those innovations, say batteries, do not keep pace with the rest, we may simultaneously observe rapid technological progress in some sectors, but slow productivity growth in the aggregate. The bottleneck created by batteries, in this case, is *endogenous* in the sense that it is the advances in the other inputs that have turned batteries into a bottleneck.

Our perspective also emphasizes how a more balanced distribution of technological progress (and research and development) can improve aggregate productivity performance. In fact, bottlenecks may signal the potential for significantly faster aggregate productivity growth in the future if and when there are major advances in areas acting as bottlenecks.

Several transformative technologies of the last three decades illustrate how bottlenecks emerge and how their alleviation can accelerate innovation and growth. High energy-density rechargeable batteries, which figuratively and literally power the mobile electronics and electric vehicle industries, provide a key example. Battery technology was a bottleneck even prior to the 1970s, when the available technology for rechargeable batteries, lead-acid electrochemistry, had low energy density, slow charging rates, a short life cycle, and an unwelcome property of releasing explosive hydrogen gas during recharging. Lead-acid batteries were succeeded in the 1970s by nickel-cadmium and nickel metal hydride (NiMH) cells, which enabled the first commercially successful gasoline-electric hybrid car, the Toyota Prius, introduced in 1997. However, the primary drive unit in the Prius remained a conventional gas engine and its NiMH battery provided only supplemental electric propulsion and regenerative breaking capacity. The battery bottleneck was substantially overcome by lithium-ion batteries, invented in 1973 and refined in the 1980s. The lithium-ion battery's high energy density not only enabled fully electric mass production vehicles, but also catalyzed a host of unforeseen innovations: a surge in onboard automotive processing power, allowing vehicle autonomy; battery-powered drone aircraft, now used in weather forecasting, emergency response, construction planning, filmmaking, and infrastructure inspection (e.g., bridge suspension systems); and the emerging electric passenger aviation industry. In awarding the prize in Chemistry in 2019 to John B Goodenough, M Stanley Whittingham and Akira Yoshino for their invention of the lithium-ion battery, the Nobel committee observed that their work had "enabled the mobile revolution."

Even more foundational to the current era is the transistor, an electronic switch capable of amplification, switching, and rectification of electrical signals (Park et al. (1976)). Through the 1950s, electromechanical switches and vacuum tubes were a clear bottleneck. Though used in virtually every kind of electronic device, telephone lines, radios, transmitters, audio amplifiers, and early computers, they were bulky, fragile and slow (Sosa (2013)). The transistor supplied a tiny, fast, and ultimately very cheap mass-produced alternative to vacuum tubes, thus break-

ing the bottleneck impeding progress in technologies as disparate as computers, long distance telephony, and audio amplifiers. Due to its extraordinary switching speed, the transistor also ushered in the age of digital communications on which all transmission of symbolic data depends. Many of the central technologies of the present—the Internet, artificial intelligence, mobile computing, digital imaging, autonomous vehicles—are transistor-dependent innovations that were largely unforeseen prior to the availability of digital switching. The transistor is estimated to be the most manufactured device in history, at 13 sextillion (10^{21}) units to date, with billions more produced each day (Iancu (2019); Laws (2018)). The transistor's immense footprint is also visible in Figure 1, where the surge in patenting in electricity and electronics and in instruments and information would not have been feasible without the transistor technology.

The Global Positioning System (GPS) constitutes a third innovation that broke a technological bottleneck and enabled a suite of technologies that have become foundational to modern life. Historically, navigating an offshore or airborne vessel required either relying on sight-lines to charted objects or applying a combination of optical instruments, precise clocks, and detailed tables. Traditional navigation was supplemented with radio-positioning systems in the 1970s, but these tools offered either poor accuracy or limited geographic coverage and hence did not penetrate beyond military and commercial shipping applications. GPS overcame these shortcomings and added a second crucial feature: atomic-level time-keeping. First launched in 1978, GPS satellites provide geolocation and date and time information to any GPS receiver on or near the earth. While GPS was built by and for the US military, the United States opened it to civilian use in 1983, after a Korean commercial airliner inadvertently navigated over Soviet airspace and was shot down. In addition to breaking the geo-positioning logjam, GPS enabled a set of highly consequential innovations that were surely not envisioned by the military planners who commissioned the system. GPS now enables precision agriculture, mining, and oil exploration; it supplies atomic time information for synchronization of power transmission systems; it enables remote surveying for geology and weather prediction; and it powers innumerable consumer-facing services such as ride-hailing, targeted advertising, and person locators.

With these motivating examples in mind, we first outline a simple conceptual framework, which helps formalize the ideas explained in the context of the three technologies discussed in the previous paragraphs. The main ingredient of our conceptual framework is that technological advances (modeled as quality improvements) in a sector depend upon simultaneous improvements in its supplier industries. Although advances in each upstream sector are potentially beneficial, these advances are complements, so that an imbalance among them is detrimental to further innovation. Our conceptual framework thus emphasizes that a balanced distribution of technological advances across sectors is important for the viability of further innovations.

This mechanism is distinct from a standard neoclassical channel where changes in input prices cause a sector to move along a fixed production possibility frontier. Our framework yields a simple estimating equation, linking growth in sectoral TFP to average TFP and dispersion (variance) of TFP among that sector's inputs. We estimate this equation using 462 four-digit manufacturing industries between 1977 and 2007 and 42 three-digit in the US between 1987 and 2007.

Our estimates indicate that greater dispersion of TFP growth among an industry's suppliers exerts a powerful negative influence on its own growth opportunities. Our preferred specification suggests that a doubling of the variance of input supplier TFP growth is associated with about 0.9 percentage points slower TFP growth for a sector.

We further document that, as conjectured, the dispersion of TFP growth among key industries has increased significantly over the last several decades. Our estimates suggest that this higher dispersion can account for essentially all of the aggregate productivity slowdown in manufacturing between the 1970s and 2007. For example, our results imply that if the cross-industry dispersion of TFP growth in manufacturing had remained at the same value as in 1977-1987, aggregate TFP growth in manufacturing would have been slightly faster in 1997-2007 than in the previous two decades.

Our methodology also clarifies which sectors are major bottlenecks and singles out a number of industries, including pharmaceutical preparation, basic inorganic chemicals, electronic connectors and surface active agents, as the leading bottlenecks. According to our results, a 20% decrease in the TFP growth of the 10 fastest growing industries with a simultaneous increase in TFP growth of each of the the bottom 50% of industries so as to keep average TFP the same would have prevented 40% of the observed rise in the variance of supplier TFP over this period and, as a result, aggregate TFP growth in manufacturing would have been 1.1 percentage points higher. Our estimates additionally reveal that surgical and medical instruments, gas engines, and industrial valves are among the most consequentially bottlenecked sectors, meaning that they are large contributors to GDP but are inhibited by high TFP growth dispersion among their suppliers.

We confirm that these empirical patterns are robust. They hold for the entire economy and within manufacturing (where TFP is better measured), and they are present in weighted and unweighted specifications, in various subperiods, with various additional controls, and with different alternative measures of productivity dispersion. We also verify that they are not driven by outliers or just by the rapid advances in computers and electronics sectors.

There is an obvious endogeneity concern in the results we present: technological trends or productivity shocks may impact supplier and customer sectors simultaneously, which could

cause us to conflate the impact of sectoral linkages with correlated shocks. As a partial remedy to this threat, we exploit international (non-US) technological opportunities as a source of variation in the variance of supplier TFP growth and obtain very similar results.

Another important concern relates to whether these results could be driven by relative price effects that change input intensity (e.g., less innovative inputs become more expensive and are used less intensively).[3] We show that this is unlikely to be the case using three complementary strategies. First, we document that our results are driven by TFP, not by quantities and prices. Second, a similar relationship holds when we focus on TFP growth dispersion among the "idea suppliers" rather than the input suppliers to an industry—that is, among sectors that an industry's innovations typically rely on for ideas. We measure these idea linkages by calculating the patent classes that a sector's own patent applications typically cite (see Acemoglu et al. (2016)). Since the citation network is distinct from the input-output network, the dispersion of TFP among the idea suppliers to an industry should be less closely linked to the relative input prices it faces. When our hypothesized mechanism is operational, however, imbalances in idea generation relevant to a sector can become another constraint on innovation, and this is what we find in the data. Third, and perhaps most directly, we document a similar relationship in patents: industries with greater patenting variance across input suppliers are less likely to patent themselves.[4]

Finally, we document analogous patterns using international data and establish that dispersion in productivity among key domestic and international supplier industries has also been a major impediment to productivity growth for several leading OECD economies.

We view our results as suggesting an important role for endogenous productivity bottlenecks and unbalanced sectoral productivity growth in constraining aggregate productivity growth. While further work is needed to test causal relationships, we can already conclude that this mechanism implies a more nuanced view of the poor productivity performance among industrialized nations over the last three decades. In particular, our hypothesis and supporting evidence suggest that once important breakthroughs arrive in sectors acting as bottlenecks, there should be an acceleration of both industry and aggregate productivity growth.

A second set of issues raised by our paper is whether the dispersion of productivity growth across sectors is inefficiently unbalanced. High dispersion may result either from changing technological opportunities or from inefficient allocation of research effort across industries.

[3]See Atalay (2017) and Baqaee and Farhi (2019) for such neoclassical effects, which arise once we depart from unitary elasticities in production.

[4]The link between patents and productivity in our sample is modest, which may be due to the imperfect correspondence between industry classifications and patent technology classes. Still, it is informative that the role of bottlenecks are detectable using both measures of innovation.

Our strategy is not geared towards answering this normative question. Positively, our evidence implies that a more balanced trajectory of technological change would generate substantial aggregate gains.

Our paper is related to a small but growing literature on the causes of the productivity slowdown. In addition to the perspectives that view the current era as either the pinnacle of very rapid but highly mismeasured productivity growth or as an exemplar of an economy running out of ideas, several other perspectives may help to explain the productivity slowdown.[5] First, and most closely related to our work, several authors have argued that productivity growth from new technologies, especially from new general purpose technological (GPT) platforms, tends to lag the underlying breakthroughs substantially because GPT-using sectors only slowly discover how to harness new technological capabilities. This idea was first proposed in the economics literature by David (1990) in the context of the effects of the electrification of American industry, which David argued took place after considerable delays. It was further elaborated by Bresnahan and Trajtenberg (1995) and Helpman and Trajtenberg (1996) who proposed mechanisms for the slow emergence of productivity gains from general purpose technologies. Closer still to our hypothesis, Brynjolfsson et al. (forthcoming) argue that productivity gains from AI and other digital technologies will trace a J-shaped curve because complementary investments and capabilities will take time to develop. Our approach, emphasizing that imbalanced innovation across sectors will act as a bottleneck, provides a specific mechanism for extensive delays in the realization of productivity gains from new technologies and platforms. Differently from these works, our paper highlights how the extent and duration of the productivity slowdown will depend not primarily on complementary investments and learning-by-doing in the GTP-using sector but on the sectoral imbalance of innovation and the speed with which breakthroughs can take place in lagging sectors.

Second, Andrews et al. (2016) provide evidence suggesting that much of the aggregate productivity slowdown is related to the poor productivity performance of non-leader firms across various sectors and countries, while leading firms have continued to experience steadily growing productivity. Several other works have emphasized specific market imperfections or failures as contributing to the productivity slowdown. These include: barriers to innovation and entrepreneurship (e.g., Decker et al. (2017), Aghion et al. (2019), and Akcigit and Ates (2019)); over-investment in automation (Acemoglu and Restrepo (2019)); insufficient government investment in R&D (Mazzucato (2013), Gruber and Johnson (2019)); and patent rent-seeking

[5]The most sophisticated version of the "running out of ideas" hypothesis is developed in Bloom et al. (2020), who argue that new innovations have become difficult in many fields, but the rate of innovation has not declined as much because the amount of effort devoted to invention and innovation has increased.

by so-called nonpracticing entities ("trolls") that discourage further innovation (Cohen et al. (2016)). Our explanation is complementary to these ideas but distinct in its focus on productivity interactions across sectors rather than primarily sector-specific or aggregate factors.

Finally, conceptually, our framework builds on models of input-output and idea linkages. Acemoglu and Azar (2020) provide a framework where innovation depends on the endogenous combinations of inputs a sector uses. Our approach here is related but emphasizes that innovation depends on how advanced (and how balanced) the set of exogenously-specified inputs is. Our framework also relates to the motivating model in Acemoglu et al. (2016), where patenting in a sector depends on the number of patents in other sectors whose patents it has typically cited. The key distinction between our approach and prior work is our focus and results on the drag that dispersion across sectors creates on aggregate innovation and productivity growth.

The rest of the paper is organized as follows. Section 2 presents a motivating conceptual framework that will guide our empirical exploration. Section 3 overviews our data sources. Section 4 presents our main results, focusing on the variance of supplier TFP growth as the measure of sectoral imbalance of innovation. This section also draws out the quantitative implications of our estimates and establishes their robustness. Section 5 provides several pieces of evidence that support our claim that the variance of supplier TFP growth captures the effects of imbalanced innovation across sectors. Section 6 presents analogous results for a cross-country panel, while Section 7 concludes. Additional information on our data, industry correspondences and robustness checks are presented in the (online) Appendix.

2 Model

In this section, we provide a motivating conceptual framework, which will then be used to derive our estimating equations.

2.1 Basic Setup

Our starting point is the idea that new product or quality innovations in a sector depend on improvements in the quality of the inputs that they use—a point emphasized by our case studies of technological bottlenecks in the Introduction. To develop this idea with minimal complexity, we consider a framework that borrows elements from existing models of input-output linkages (such as Long and Plosser (1983), Acemoglu et al. (2012), and especially, Acemoglu and Azar (2020)) and from canonical quality-ladder models (such as Aghion and Howitt (1992) and Grossman and Helpman (1991)).

Suppose in particular that there are N sectors, denoted by $i = 1, 2, \ldots, N$. Assume also that the production function of sector i at time t is

$$Y_{it} = B_i A_{it} L_{it}^{1-\sum_{j\in S_i}\alpha_{ij}} \prod_{j\in S_i} X_{ijt}^{\alpha_{ij}}. \tag{1}$$

Here, Y_{it} denotes the output of sector i at time t, A_{it} is the productivity of this sector at time t, and B_i is a normalizing constant.[6] In addition, each sector uses labor, L_{it}, and inputs that are necessary for production, which are those in the time-invariant set S_i.[7] For simplicity, these inputs are assumed to be combined with a constant returns to scale Cobb-Douglas technology, where α_{ij} are input shares and $1 - \sum_{j\in S_i}\alpha_{ij}$ is the share of laborin production.

The most important building block of our conceptual framework is the way in which quality improvements (or equivalently the introduction of new, higher-quality products) take place. Here, we introduce the quality-ladder structure by assuming that $A_{jt} = \lambda^{n_{jt}}$ where $\lambda > 1$ and n_{jt} is the number of innovations this sector has experienced in the past. Each innovation, therefore, increases productivity by a factor λ.

Our critical assumption is that the arrival rate of innovation depends on the distribution of technologies of all the inputs the sector uses:

$$\phi_{it} = H\left(\sum_{j\in S_i}\alpha_{ij} h(A_{jt})\right), \tag{2}$$

where ϕ_{it} denotes the arrival rate of a new innovation at time t, and h and H are monotone continuous functions, and we normalize $H(0) = 0$.[8] Different choices for these functions give different relationships between the quality distribution of the inputs of the sector and its innovation propensity. For example, we could take $h(x) = x^{\rho}$ and $H(x) = x^{1/\rho}$ to obtain a CES (constant elasticity of substitution) aggregator. Particularly important in this context is whether the function h in equation (2) is convex or concave. The former corresponds to a situation in which innovation in each sector is determined by how advanced its most advanced inputs are, which makes innovation across different sectors substitutes and implies that greater (mean-preserving) dispersion in innovations across inputs helps innovation. Alternatively, the

[6] $B_i = \left((1-\sum_{j\in S_i}\alpha_{ij})^{1-\sum_{j\in S_i}\alpha_{ij}} \prod_{j\in S_i}\alpha_{ij}^{\alpha_{ij}}\right)^{-1}$. See Acemoglu and Azar (2020) for more details on this functional form.

[7] It is straightforward to allow these sets to be time-varying, but we do not do so to reduce notation. In our empirical work, we explore models both with and without time-varying input sets.

[8] We have equated the importance of an input for innovation to its share in production, α_{ij}. This is not necessary for any of our main arguments, but it is the benchmark functional form assumption we use in our empirical work. We also consider an alternative where the importance of an input innovation is measured by the number of patent citations between industries. See Acemoglu and Azar (2020) for a discussion.

concave case, arises when innovations across different inputs are complements, so that greater (mean-preserving) dispersion hinders innovation. We find this concave case to be empirically more relevant, because it captures the intuitive idea, already highlighted in our case studies in the Introduction, that new product and quality improvements necessitate simultaneous improvements in a range of inputs, and if some of the relevant inputs fall behind, they will act as a bottleneck, retarding technological progress.[9] Note also that, in both the convex and the concave case, because both h and H are monotone, a higher level of technology for any input always helps innovation in the sector in question.

A second-order Taylor expansion of the right-hand side of equation (2) around its mean gives:

$$\phi_{it} \approx H\left[\alpha_{ij} h(\bar{A}_{it}) + h''(\bar{A}_{it}) var(\{\alpha_{ij} A_{jt}\}_{j \in S_i})\right],$$

where $\bar{A}_{it} \equiv \sum_{j \in S_i} \alpha_{ij} A_{jt}$ is the (cost-share weighted) mean of productivities of the inputs to sector i, and $var(\{\alpha_{ij} A_{jt}\}_{j \in S_i})$ is the (weighted) variance of the productivity as of the inputs to sector i. Next, taking a first-order expansion of H around 0 and also approximating $h(\bar{A}_{it})$ by $h'(\bar{A}_{it})\bar{A}_{it}$ gives:

$$\phi_{it} \approx \eta^i_{mean} \bar{A}_{it} + \eta^i_{variance} var(\{\alpha_{ij} A_{jt}\}_{j \in S_i}), \tag{3}$$

where $\eta^i_{mean} \equiv H'(0)h'(\bar{A}_{it})$ represents the effect of the mean of the technological advances across inputs, which we always control for in our empirical work, while $\eta^i_{variance} \equiv H'(0)h''(\bar{A}_{it})$ captures the effect of dispersion across inputs (holding the mean constant). Equation (3) will be the basis of our empirical work in the rest of the paper. The estimates of the parameter $\eta_{variance}$ will show whether, in terms of our framework, the function h is convex or concave. This coefficient will also indicate the extent to which unbalanced innovations across key inputs in the economy may hold down aggregate productivity growth.[10]

To illustrate this point succinctly, suppose that $S_i = S$ for all i and some $S \subset \{1, \ldots, N\}$ and $\alpha_{ij} = \alpha_j$ for all i and $j \in S$. Suppose also that h is concave, so that $\eta_{variance} \equiv H'(0)h''(\bar{A}_t) < 0$, and we start with $A_{jt} = \bar{A}_t$ for all $j \in S$. Then consider a mean-preserving spread of the A_{jt}'s so that they now have a weighted variance, $var(\{\alpha_{ij} A_{jt}\}_{j \in S_i})$, given by σ^2. Equation (3) implies

[9] The inputs that need to make technological advances for sector i to innovate successfully may be a subset of the inputs in S_i, but since we do not have a way of determining which subset of inputs is important for innovation in our empirical work, we assume that all inputs in S_i are relevant, and then verify robustness using other measures of industry linkages.

[10] Even when $\eta^i_{variance} < 0$, an increase in the productivity of an input supplier industry is always beneficial (and thus this effect is less than the impact through the mean, η^i_{mean}), because the functions h and H are monotone.

that the aggregate productivity of the economy will be reduced by $\eta_{variance}\sigma^2$. So if σ^2 and $\eta_{variance}$ are both large, there will be a sizable negative impact on aggregate productivity.[11]

2.2 Endogenous Innovation Effort

It is straightforward to endogenize innovation and characterize the general equilibrium.[12] While the potential endogeneity of innovation does not play an important role in our empirical work, it is nevertheless useful to consider it to motivate our later discussion of potential inefficiencies from unbalanced innovative effort. We add this channel to the model by modifying equation (2) to

$$\phi_{it} = \frac{1}{\gamma} H \left(\sum_{j \in S_i} \alpha_{ij} h(A_{jt}) \right)^{1-\gamma} z_{it}^{\gamma}, \tag{4}$$

where $\gamma \in (0,1)$ and z_{it} is research effort devoted to innovation in industry i at time t (e.g., overall research spending or research-related resource use such as scientific effort). This specification implies that there are intra-temporal diminishing returns to research effort in a given field, which could arise from crowding out when multiple researchers pursue similar ideas simultaneously. We include $1/\gamma$ as a constant in front of the H function for simplicity.

Suppose also that the per-unit cost of research in industry i is κ_i, and the reward to an innovation in the sector at time t is π_{it}. The cost κ_i depends on the opportunity cost of research-related resources in non-research activities, but may additionally include sector-specific distortions, as well as misperceptions or fads among researchers (e.g., researchers being motivated to pursue a particular field beyond its social value). We interpret the reward π_{it} as a market-outcome determined by prices, market sizes and markups, though fads and misperceptions may affect rewards as well.

Given this setup, privately optimal research effort devoted to sector i at time t will be

$$z_{it}^{*} = \left(\frac{\pi_{it}}{\kappa_i} \right)^{\frac{1}{1-\gamma}} H \left(\sum_{j \in S_i} \alpha_{ij} h(A_{jt}) \right),$$

[11] As we explain in the next subsection, it may not be possible to reduce the dispersion of technological progress across sectors without affecting the mean. In particular, such a mean-preserving dispersion reduction would require that the cost of improving technology in any sector is the same.

[12] The general equilibrium will additionally involve the determination of the equilibrium wage rate (and the allocation of labor across sectors) and the equilibrium interest rate (as a function of the aggregate growth rate of the economy). We do not derive these additional (and standard) aspects of the general equilibrium.

and thus

$$\phi_{it}^{*} = \frac{1}{\gamma}\left(\frac{\pi_{it}}{\kappa_i}\right)^{\frac{1}{1-\gamma}} H\left(\sum_{j\in S_i}\alpha_{ij}h(A_{jt})\right), \tag{5}$$

which is proportional to the exogenous-specified success probability in equation (2). This ensures that equation (3) applies as before, and highlights that whether there is endogenous or exogenous innovation success probabilities is not central for our empirical work.

Equation (5) emphasizes that, to the extent that the κ_i's vary across sectors for reasons unrelated to the social cost of innovation in sector i, the unequal (unbalanced) rates of technological progress across sectors could be inefficient. In such a scenario, policies that reduce the dispersion of the rates of technological progress across sectors may improve the allocation of resources. For example, if the marginal cost of innovation were the same across sectors, a social planner could reduce dispersion without affecting mean technology, improving aggregate productivity (and welfare). Conversely, if differences in κ_i's across sectors reflect differences in the social costs of innovation, then it may be infeasible to reduce the sectoral dispersion of technological progress without lowering mean productivity in the economy. Since we do not know where differences in the rate of innovation across sectors come from, these observations caution against drawing strong normative conclusions from the results that follow.

3 Data Sources

The data sources that form the backbone of our paper combine time-series for industry TFP growth with input-output linkage data. For manufacturing industries, we use data from the NBER-CES Manufacturing database. These data are sourced from the Annual Survey of Manufacturers and include annual industry-level data from 1958-2011 on output, employment, input costs, investment, capital stocks, TFP, and industry-specific price industries. We include 462 manufacturing industries corresponding to six-digit NAICS codes. As is standard throughout the literature, TFP is defined as the residual of the change in real output after subtracting the cost-share weighted change in each of five factors (capital, production labor, non-production labor, energy, and non-energy materials). We supplement the manufacturing data with annual TFP estimates for non-manufacturing industries from the BLS Multifactor Productivity Database from 1987-2011. As with the manufacturing data, TFP outside manufacturing is defined as the difference between real output growth and a shares-weighted combination of growth in five inputs (capital, labor, energy, materials and purchased services).[13]

[13]The BLS also produces similar statistics for aggregated three-digit manufacturing industries. Although we do not use these BLS measures in our analysis, these statistics are highly correlated with the Census data multifactor productivity measures for manufacturing

We construct input-output tables using the detailed Make and Use Tables provided by the US Bureau of Economic Analysis from 1977-2002, which are available every five years, corresponding to the years of the Economic Census. These tables provide information on the amount that each industry spends on various commodities and the amount that they produce of each commodity, respectively. From these two tables, we construct the dollar value that each industry i uses as intermediates from other industries. Since each year's release of these tables is in industry codes particular to that year's classification, we convert each table to a set of time-consistent NAICS-based industry codes, the details of which are documented further in the Appendix. Table 1 presents summary statistics for these main economic variables, both across all manufacturing industries and the distribution of them across supplying industries. Panel A shows the results for manufacturing industries only, while Panel B depicts averages for all sectors. We see that average TFP growth over a five-year period across manufacturing sectors over this time period was 1.8 percentage points. The average TFP growth of upstream industries is substantially higher, reflecting the fact that more productive industries are used more intensely as intermediate inputs.

As an alternate measure of productivity by sector, we also collected information on patenting by industry from 1976-2006 from the NBER patent data project. Each patent in this dataset is assigned to a US Patent code, which we convert to a SIC-based industry classification using the mapping produced by Kerr (2008). We use this dataset to construct a time-series for total patents and total patent citations for each industry. We supplement patent citations data with information on "high-value" patenting by industry constructed in Kogan et al. (2017), who identify these patents from stock-market responses to patent announcements as those that are in the top quintile of the distribution of stock market excess returns. As with the other patenting measures, we assign high-value patents to industries using the mapping from USPTO codes to SIC-based industry codes. Panel A of Table 1 shows the growth in these patenting measures, both on average and among supplying industries. Consistent with Figure 1, there is an enormous growth in patenting over this time period, while the input-adjusted growth in patents is slightly lower than the simple average across manufacturing industries.

As an additional measure of innovation, we follow Acemoglu et al. (2016) and construct citation-based innovation linkages between industries. Specifically, we calculate the fraction of citations for patents in industry i that come from patents in industry j. In order to simultaneously capture potential time-variation in this network of citations and account for the lags in citation patterns, we construct the citation network for year t using all citations occurring between years $t-12$ and t to patents granted during years $t-8$ through $t-12$. Thus, the 1987 citation network is calculated as the total number of citations on patents that were granted be-

tween 1975 and 1979, the 1992 citation network is calculated using citations to patents granted between 1980 and 1985, and so on. We also show robustness of our results to using a single time-invariant citation network that pools all years.

Lastly, we supplement the domestic US data with data for select European countries. We take data on value-added and TFP from the 2012 EU KLEMS Growth and Productivity accounts. In this exercise, we use data from 1982-2007 for 30 industries in Spain, France, Austria, Finland, the Netherlands, Italy, Germany and the UK. We combine these data with country-specific input-output tables from the World Input Output Tables for 2000, the earliest year available. The relevant entry in the world IO table, $\alpha_{ik,jl}$, is the share of inputs from industry i in country k that came from industry j in country l.[14] Panel C of Table 1 shows the TFP growth among this sample. Overall, the growth in TFP across each five-year period in this sample was 4.6 percentage points.

4 Sectoral Imbalances and Productivity Growth

This section presents our main results linking the total factor productivity (TFP) growth of an industry to the dispersion of productivity growth among its suppliers—corresponding to imbalance, under our hypothesis. Concretely, we estimate a version of equation (3), derived above, using data on 462 four-digit manufacturing industries between 1977 and 2007, and 42 non-manufacturing industries in 1987-2007. We also report the quantitative implications of these estimates and document their robustness to control vectors, sample periods, and sources of variation in productivity growth.

4.1 Main Results

Our main estimating equation is the empirical analogue of (3):

$$\begin{aligned} \Delta TFP_{it} &= \beta_{mean} \sum_j \alpha_{ijt-1} \Delta TFP_{jt} \\ &\quad + \beta_{variance} \text{VAR}\left(\Delta TFP_{jt}\right) + \mathbf{X}'_{it-1}\beta_{other} + \delta_t + \varepsilon_{it} \end{aligned} \tag{6}$$

[14] In our baseline specification, this share includes inputs from all other countries. We explore alternative definitions and show the robustness of the results to using only the domestic input-output table specific to each country.

where t refers to five-year time periods, ΔTFP_{it} is the TFP growth of industry i during the five-year time interval denoted by t,

$$\text{VAR}\left(\Delta TFP_{jt}\right) \equiv \sum_j \alpha_{ijt-1} \left(\Delta TFP_{jt} - \sum_j \alpha_{ijt-1} \Delta TFP_{jt} \right)^2,$$

and $\sum_j \alpha_{ijt-1} \Delta TFP_{jt}$ is the average TFP growth among the suppliers of industry i in the five-year time period, calculated using the share of inputs from the industry in question in the intermediate costs of industry i in the beginning of the time period (α_{ijt-1}). The variance of TFP growth among the suppliers of industry i is also computed using cost shares as weights. In addition, $\mathbf{X}'_{it-1}$ denotes a vector of other (predetermined) covariates, which include sector fixed effects (allowing sector-specific linear trends) in some of our specifications; δ_t denotes a full set of time dummies; and ε_{it} is a heteroscedastic and potentially serially-correlated error term, capturing all omitted factors.

This equation is comparable to our model-derived equation (3) above, with several operational refinements. First, we use TFP growth as our primary measure of innovation since we do not have direct measures (though we will look at patenting as well). Second, instead of relating innovation to the *level* of technologies across inputs as in equation (3), we link TFP *growth* in each sector to the TFP growth rate across inputs, since the level of TFP is not well defined. Third, we have included an error term and additional covariates. Fourth, instead of the sector specific coefficients in front of the mean and the variance in equation (3), the η^i_{mean} and $\eta^i_{variance}$'s, we have imposed constant coefficients, to be interpreted as local average treatment effects.

Throughout, even though we always control for the mean effect of supplier TFP growth, the main coefficient of interest for our study is $\beta_{variance}$, which captures the effect of supplier TFP growth/innovation dispersion on a sector's productivity/innovation, holding the mean of supplier TFP growth/innovation constant. We expect this coefficient to be significantly negative if, as we hypothesize, imbalances in the rate of technological progress across suppliers to an industry imposes a productivity penalty on that industry.

Table 2 reports estimates of equation (6) for 462 four-digit manufacturing and 42 three-digit non-manufacturing industries. Panel A is for manufacturing industries, for which TFP estimates are more reliable and available for a longer time period. Panel B combines the manufacturing and the non-manufacturing industries to include the full set of sectors. Odd-numbered columns include no covariates other than time dummies, while even-numbered columns (and column 9) also include industry fixed effects, thus allowing each industry to have its own linear time trend in TFP. The standard errors account for arbitrary heteroscedasticity and serial correlation at

the industry level throughout. Our baseline regressions shown in columns 1 through 6 are unweighted. We present results weighting industries by their share of 1987 real value added in columns 7 through 9.

Column 1 shows the relationship between industry TFP growth and mean supplier (upstream) TFP growth—focusing only on the first term in equation (6). We detect a strongly positive relationship between mean supplier TFP growth and downstream industry TFP growth. Adding the variance term in columns 2 and 3 strengthens the effect of mean supplier TFP growth, and more importantly, shows a precisely-estimated and quantitatively large *negative* relationship between the variance of supplier TFP growth and industry TFP growth. For example, in Panel A the coefficient estimate of the variance term is -0.744 (standard error = 0.121) in the base model in column 2. Adding linear industry trends modestly increases this coefficient to -0.912 (standard error = 0.118). When we include non-manufacturing industries in Panel B, the point estimates are similar, even if a little larger. Figure 3 depicts the industry-level variation producing these estimates. Specifically, we report binscatters for the regression model in column 2 of Panel A. The left panel depicts the strong positive relationship between average supplier TFP growth and industry TFP growth, and the right panel showcases the strong negative relationship between the variance of supplier TFP growth and industry TFP growth.

The specification in (3) is a natural one, capturing all the effects from supplier TFP growth dispersion by the variance term (which, as we have shown, can be thought of as a second-order approximation to any type of complementarities between innovations and inputs). Nevertheless, it is useful to see whether both very well performing and very poorly performing supplying sectors impact TFP growth. To investigate this question, columns 4 and 5 replace the variance term with TFP growth in the 10th and 90th percentiles of the (weighted) TFP distribution of suppliers (while we still continue to control for meat supplier TFP growth). Consistent with our hypothesis, bottom decile supplier TFP growth predicts *faster* own-industry TFP growth, while the top decile supplier TFP growth predicts *slower* own-industry TFP growth (with both relationships typically statistically significant at 5% or less). Lastly, columns 7 through 9 replicate the main specifications weighting each industry by its share of real-value added in 1987. These weighted estimates are very similar to the unweighted specifications.

Overall, the estimates in Table 2 uniformly show a negative estimate of TFP growth dispersion across a sector's suppliers. In terms of our motivating conceptual framework, this suggests that productivity growth and innovation in a sector is held back when advances among its suppliers is imbalanced. In the rest of the paper, we explore the quantitative implications of our estimates, demonstrate their robustness, and provide additional evidence that they capture the effects of imbalanced innovations across suppliers.

4.2 Quantitative Implications

The Table 2 estimates imply that an imbalance in productivity growth across sectors could be a drag on aggregate growth. Temporarily deferring robustness checks, we explore whether such sectoral imbalances could be a quantitatively meaningful contributor to the productivity slowdown in the US. For this to be the case, two conditions need to be satisfied. First, the coefficient estimates in Table 2 must be economically large. Second, the dispersion of sectoral TFP growth must have increased over the decades during which we witnessed the productivity slowdown.

Figure 4 confronts the latter issue by plotting the evolution of the variance of TFP across manufacturing industries. The top two plots in Figure 4 depicts the simple variance of TFP growth across manufacturing industries, while the middle and bottom plots show the average variance of the industry supplier TFP growth in manufacturing and in the economy overall, respectively. Within manufacturing and in the economy overall, there was a striking rise in the dispersion of sectoral productivity growth in the US economy over the last several decades. This is true both overall and when weighting industries by their input share. Quantitatively, the TFP variance in manufacturing was about 0.002 before the 1970 and now is three times as large, around 0.006. As suggested by the patenting time-series in Figure 1, a large portion of this increase in TFP variance through the 1990s is accounted for by the electronics and computer sectors. The right panels of Figure 4 documents that when these sectors are taken out, the rise in the variance of TFP growth is significantly smaller, though still present in the recent period. When we zoom out to include non-manufacturing industries, there is again a large increase in the variance of TFP growth from 1987-92 to the present, but the pattern is not monotone, perhaps reflecting the fact that TFP is measured less reliably outside of manufacturing.

How much of the productivity slowdown can the rising variance of TFP growth potentially explain? Figure 5 addresses this question by applying the value-added weighted estimates reported in column 7 of Panel A and Panel B in Table 2. We find a sizable productivity penalty from TFP growth dispersion. The estimates imply that TFP dispersion reduced TFP growth by over 0.5 percentage points in manufacturing in each of the recent 10-year periods (as shown by the green bars) and by about 0.2 percentage points in the entire economy in each five-year period. In combination with the trends in TFP variance reported Figure 4, we find that rising TFP variance accounts for a substantial part of the productivity slowdown. If, counterfactually, TFP variance in manufacturing from 1997-2007 had remained at its average level from 1977 through 1987, our estimates imply that aggregate TFP growth in manufacturing between 1997-2007 would have been 3.1% (as shown in the orange bars) instead of the observed TFP growth of

-1.6 (as shown in the blue bars). This counterfactual growth rate is slightly more rapid than the 3% actual TFP growth in manufacturing in the previous two decades. The implied effects for the overall economy in the right panel are also sizable: had the overall variance of TFP growth stayed at its 1987 through 1992 level, aggregate TFP growth would have continued to rise over the sample period, averaging 1.7 percentage points higher than its observed growth rate in each five-year period. Thus, the quantitatively sizable estimates in Table 2 can potentially account for the bulk of the US productivity slowdown in recent decades.

To provide more detailed insight into these aggregate relationships, we explored the identities of the sectors that have contributed to this quantitative effect. The variance of supplier TFP in manufacturing increased over this period both because lagging industries failed to grow and because leading industries pulled away from the rest. Panel A of Table 3 lists illustrative examples of the fast-growing industries that have had the largest effect on the variance between 1997 and 2007.[15] These include electronic computers, computer storage devices and semiconductors. To gauge the economic leverage of these outlier industries, consider a hypothetical mean-preserving contraction of TFP growth dispersion: reduce TFP growth of the 10 fastest growing industries between 1997 and 2007 by 0.2 percentage points and increase the TFP growth of each of the the bottom 50% of industries just enough to keep the average TFP growth constant. In this scenario, the variance of supplier TFP growth between 1977 and 2007 would have been 40% lower and aggregate TFP growth in manufacturing would have been 1.1 percentage points higher.

The remaining panels of Table 3 round out the evidence on bottleneck industries. Panel B reports illustrative examples of slow-growing industries that became the biggest bottlenecks over the same time periods. These include pharmaceutical preparation, basic organic chemicals, printed circuit assembly and turbine generators. Panel C reports example industries that are *bottlenecked*—that have been most held back by the uneven innovation across their suppliers. These include surgical and medical instruments, gas engines, and industrial valves.

4.3 Robustness

Table 4 investigates the robustness of our results to the inclusion of a battery of controls and different specifications. For brevity, we focus on the manufacturing sample and report analogous results for the full sample in the Appendix (see Table A5). Panel A documents robustness for the specification without industry fixed effects, Panel B includes industry fixed effects that allow for linear trends in industry TFP, and Panel C shows robustness for the weighted regressions,

[15] Appendix Table A4 reports the full set of top ten industries for all panels of Table 3.

which may be more informative for aggregate magnitudes.

We begin in column 1 of Table 4 by reporting for reference our baseline specification, which includes only the mean and the variance of TFP (as well as time dummies). In column 2, we estimate the same models but now using 10-year periods rather than the five-year intervals in Table 2. This specification purges higher-frequency variation in TFP and focuses on the downstream impacts of innovations over a longer time interval. The results from these models are similar to the baseline estimates.

Our estimating equation (3) defines sectors that are falling behind as those that have relatively slow TFP growth in the contemporaneous five-year period. However, if high variance in the current period reflects mean reversion from rapid growth in the recent past, this would not correspond to an imbalance but rather a potential rebalancing. Column 3 provides a check on this possibility by including the covariance between supplier TFP growth in the current and prior periods in our regressions.[16] Intuitively, this covariance term accounts for the potential role of persistence (or its opposite) in industry-level TFP changes. We find that accounting for the covariance of TFP across periods does not change the relationship of primary interest: the coefficient on the variance term in Panel A is now only slightly larger, -0.640 (with a standard error of 0.127), while the covariance term is relatively small and imprecisely estimated. The estimate on the covariance term is larger and statistically significant in Panel B, but in both Panels B and C the coefficient on the variance of upstream TFP growth is largely unaffected by the inclusion of the covariance term. We infer from these results that the first and second moments of the upstream TFP distribution provide informative measures of sectoral imbalances.

The subsequent columns of Table 4 provide additional robustness checks. Column 4 adds a lag of the dependent variable to test whether mean reversion in industry TFP confounds the estimates. The lagged dependent variable is insignificant and small in Panels A and C, and small and opposite-signed, albeit statistically significant, in Panel B. Accounting for mean reversion does not substantially change our estimates of the effects of supplier TFP dispersion: the coefficient estimate of the variance term is -0.747 (standard error = 0.115) in Panel A and -0.908 (standard error = 0.120) in Panel B.

Another factor that could affect measured industry TFP is changing import penetration. Columns 5 controls for average imports from China (from Autor et al. (2013)). Accounting for imports does not appreciably change the coefficient on supplier TFP variance.

[16]Specifically, we calculate the covariance among suppliers between the TFP growth in the previous five-year period ($t-10$ to $t-5$) and the curent period ($t-5$ to t), weighting each supplying industry by their input share in $t-10$.

We noted the importance of the electronics and computer sectors above. Column 6 shows that the negative relationship between industry TFP growth and supplier TFP dispersion holds even when computers and electronics manufacturing (NAICS code 334) are excluded from both the calculation of the upstream metrics and from the estimation sample. The variance term is now less precisely estimated, as expected, with these key sectors excluded. Nevertheless, it remains statistically significant at 5% or less in all of our specifications: -1.231 (standard error = 0.587) in Panel A, -1.197 (standard error = 0.624) in Panel B, and -1.770 (standard error = 0.869) in Panel C. These estimates highlight that our hypothesized mechanism is present even when the ICT and electronics sectors are excluded. Simultaneously, they also highlight that the ICT and electronics sectors showcase our mechanism and contribute substantially to its identification and quantitative implications (as indicated by the examples in Table 3).

Column 7 shows a similar relationship to our baseline estimates when we use a fixed input-output matrix, rather than the time-varying input-output matrix from our baseline specification. Column 8 probes robustness to our definition of input shares. Here, we define upstream shares α_{ijt-1} as total cost shares rather than as shares among the cost of intermediates as in our baseline specification. These two share measures will differ to the extent that the intermediate share varies across sectors. The results are once again very similar.

Finally, column 9 provides a placebo check by including the mean and variance of future input TFP growth rather than current TFP growth. Under our hypothesized mechanism, the mean and variance of future TFP growth should have no impact on productivity today. Reassuringly, our estimates confirm this expectation, and future input TFP growth and variance have very little predictive power for an industry's current TFP. Results are again similar if we also include the contemporaneous mean and variance of TFP growth (not reported).[17]

Since our empirical analysis is confined to 426 manufacturing or 504 total industries, our estimates will not capture any imbalances in innovation or productivity growth happening at more disaggregated levels. To explore whether these more micro imbalances may also matter and further probe the robustness of our results, in the Appendix we use estimates of within-industry, across-establishment TFP growth the U.S. Census Bureau's Dispersion Statistics on Productivity (DiSP). Figure A2 shows that these measures of dispersion have also increased during our sample period. Nevertheless, Appendix Table A3 documents that the average up-

[17] Appendix Table A1 confirms that our results are very similar when we drop outliers or use an outlier robust estimator that minimizes the effects of individual observations on coefficient estimates. Appendix Table A2 shows results using industry value added per worker instead of TFP (Panel A) and estimates with lagged mean and variance of supplier TFP (Panel B). The results are broadly similar to those in Table 2, but more sensitive to outliers in Panel A and significantly weaker in Panel B, perhaps reflecting the fact that TFP variance is very weakly correlated over time. (The correlation between lagged TFP growth and current TFP growth is 0.05.)

stream TFP growth dispersion among input suppliers, when added to our regression, is itself not statistically significant and does not change the relationship between our measure of supplier TFP growth dispersion and own TFP growth. A more systematic investigation of these issues necessitates firm-level data, which we leave for future work.

In summary, these estimates, and those in the Appendix Table A5 for the entire economy, confirm that the negative relationship between industry TFP growth and supplier TFP variance is statistically significant, pervasive, and largely unaffected by the inclusion of a variety of potential confounders.

4.4 Instrumental-Variables Estimates

Since the estimates above regress industry-level TFP growth on the contemporaneous TFP growth of its suppliers, productivity shocks that are common across several industries might generate mechanical correlations between our right-hand side and left-hand side variables. Isolating industry productivity changes that emanate from common technological developments across several advanced economies can alleviate this simultaneity concern. This is what we do in Table 5 by exploiting changes in industry TFP in major OECD countries, as reported by the 2012 EU KLEMS Growth and Productivity accounts. In Panel A, we use the mean and variance of supplier TFP in Germany, France and the UK as instruments for the corresponding variables in the United States. To purge measurement error in these instruments, Panels B and C use the analogous measures computed from the ranks of TFP growth by industry within country.[18]

The first-stage F-statistics are given at the bottom of Panel A and Panel B of Table 5, and are somewhat low in Panel A but higher in Panel B, where we use the rank of TFP as the instrument (the first stages are reported in Table A6 in the Appendix). Nevertheless, the Limited Information Maximum Likelihood estimates in Panel C, which are consistent even in the presence of weak instruments, confirm that our estimates are not driven by weak instrument problems.

The instrumental variables estimates of the relationship between supplier TFP mean and variance and industry TFP growth correspond closely to our earlier OLS estimates. This can be seen by comparing columns 1 and 2 of each panel of Table 5, which report OLS estimates for the IV subsample, and columns 3 and 4 of each panel, which report the corresponding IV

[18]We calculate these instruments using the US-based IO table but taking the TFP growth across industries from each of these three European countries. Because the international industry data are more aggregated than our underlying NAICS data, TFP growth in six-digit NAICS industries that map to the same international industry code is assigned the same value of the instrument.

models. In columns 1 and 3, which do not include industry fixed effects, the OLS coefficient on the variance term is -0.907 (standard error= 0.166), while the IV estimate for the same variable in the same specification is -0.902 (standard error = 0.379) using the level of TFP growth and is -0.667 (standard error =0.326) using the rank of TFP as the instrument. Adding industry fixed effects (even-numbered columns) reduces precision, as expected, but the point estimates remain stable across specifications.[19]

Lastly, since TFP growth in other countries could also affect US TFP growth through its effect on imports, columns 5 and 6 show IV estimates that control for the growth in imports to the US from each country (France, Germany and the UK) in the focal industry and in its suppliers. This makes little difference in practice, suggesting that this channel is not an important source of bias in our estimates.

The congruence between the OLS estimates and those exploiting TFP changes in other leading economies bolsters our confidence that our main results are not driven by shocks that are common to industries and their suppliers in the US. We also note that since the IV coefficient estimates are similar to the OLS estimates from Table 2, the implied quantitative magnitudes are comparable as well.

4.5 Prices, Quantities and Productivity

Could these patterns be explained by mismeasurement of TFP? In a standard neoclassical setting, industries benefit when their suppliers increase their productivity because this reduces input costs (e.g., Acemoglu et al. (2012)). If TFP were measured correctly, it would be unaffected by fluctuations in employment, demand factors, and input costs that induce industries to move along their production possibility frontiers rather than changing those frontiers. If TFP were mismeasured, however, these neoclassical effects could erroneously spill over to TFP estimates, and neoclassical relative price effects that arise with non-unitary elasticities, explored in Atalay (2017) and Baqaee and Farhi (2019), might confound our results.

We investigate the role for these neoclassical channels in Table 6, where we add the mean and variance of supplier prices and employment levels to our baseline regressions. Panel A of the table uses the time-varying input-output linkages, while Panel B focuses on the time-invariant input-output network case. For comparison, columns 1 and 2 replicate our baseline estimates.

The various specifications reported in Table 6 indicate that controlling for these neoclassical channels does not qualitatively change the relationship between supplier TFP variance and industry TFP growth (and these additional variables themselves are typically insignificant).

[19]See Appendix Table A7 for robustness of these IV results to using only a subset of countries as instruments.

For example, when we include the mean and variance of supplier *prices* in column 3 of Panel A (without industry fixed effects), the TFP variance term has a coefficient of -0.686 (standard error = 0.232), which is 90% as large as the baseline estimate, though less precisely estimated. In Panel B, the estimate of the variance term is actually larger than in the baseline model, -1.061 (standard error = 0.237). When we include the mean and variance of supplier *employment levels* in column 5 (again without industry fixed effects), the coefficient estimate on TFP variance is -0.703 (standard error = 0.118) in Panel A and -0.639 (standard error = 0.161) in Panel B, both nearly identical to the baseline estimates in column 1. The results are also similar when we include both sets of variables (prices and employment) together. When we additionally include industry fixed effects, the estimates are again very similar to our baseline results.

The evidence in Table 6 suggests that the relationship between supplier TFP and industry TFP detected above is not a reflection of (potentially mismeasured) neoclassical effects. Instead, we believe that it captures economic effects working through the innovation or product quality mechanism identified by our model. We next offer more direct evidence on this mechanism.

5 Innovation

In this section, we investigate whether innovation, as encoded in patents, is one of the underlying mechanisms accounting for the results presented so far. Tables 7 and 8 provide evidence on the role of innovation linkages in productivity growth. In Table 7, we estimate models akin to our baseline setup for TFP, but replacing the input-output network with the patent citation network calculated in Acemoglu et al. (2016).[20] The regressions in this table thus assess whether greater dispersion in TFP growth among industries typically cited by the patents of a sector is associated with that sector itself experiencing slower TFP growth. Since neoclassical effects working through prices and quantities should be less directly relevant for industry linkages based on knowledge flows, this exercise also allows us to explore whether, as hypothesized, the estimated adverse impact of TFP dispersion stems in part from its relationship to innovation.[21]

The results in Table 7 document a powerful negative relationship between the variance of TFP growth in an industry's citation network and its own TFP growth. The first four columns show unweighted regressions, while columns 5 through 8 present estimates weighting

[20] As described in Section 3, in order to balance the lagged timing of citations with the noise inherent in measuring the citation network, we define the 1987 citation network to include the set of patents registered between 1975 and 1980, regardless of when in the following years they were cited, and similarly for later years. We find that results are analogous when using a time-invariant citation network that pools all years. See Panel B of Table 7.

[21] However, the moderately high correlation between the citation-based and the input-output-based mean and variance measures implies that separating these two effects is difficult in practice.

each industry by its share of value added in 1987. In column 1 of Panel A, which does not include industry trends, the variance term has a coefficient of -1.201 (standard error = 0.388). The coefficient of interest is larger and still highly significant, albeit less precise, when we additionally include industry fixed effects in column 2. The point estimates are even larger in the corresponding weighted specifications in columns 5 and 6.

Columns 3-4 and 7-8 present a horse-race between industry linkages based on citation networks and input-output networks. Across all columns, the estimates pertaining to the input-output network are more precise, confirming the robustness of our baseline results. The variance of TFP growth in the citation network remains negative across all columns, but is no longer precisely estimated in these horse-race specifications. This likely reflects the greater mismeasurement of idea linkages via citations than our baseline linkages based on inputs. On balance, we interpret the estimates as showing that both sets of network linkages appear to simultaneously exert leverage on TFP growth in downstream sectors.

To probe robustness, Panel B and Panel C of Table 7 fix the industry linkages (both citations and input-output shares) across industries at their 1987 level, or alternatively use a more aggregated industry-level citation network. Both sources of linkages point in the same direction as Panel A and confirm that greater dispersion of TFP growth among linked industries predicts slower TFP growth.

Our evidence so far establishes a robust relationship between an industry's TFP growth and the TFP growth of its suppliers. Because TFP is an inferred rather than observed measure of technological progress, these relationships constitute indirect evidence linking supplier industry innovations to downstream customer industry innovations. Table 8 supplements this evidence by testing for a direct link between an industry's innovative output and the distribution of innovations among its suppliers. Specifically, we estimate a variant of equation (6) where we replace industry TFP with industry patenting on the left-hand side, and use the mean and variance of the suppliers' patent growth on the right-hand side. We focus on three measures of supplier patents: the number of patents (Panel A); the citation-weighted count of patents (Panel B); and the count of patents in the top 20% of value (Panel C) as defined following Kogan et al. (2017). We view the latter two measures as most reliable since they capture patent quality as well as quantity. The sample now runs from 1982 to 2002, and we conduct the analysis using 490 SIC-based industry codes in manufacturing.[22]

The results in Table 8 illustrate the strong linkage between an industry's innovations and its suppliers' innovations, and confirm that sectoral imbalance in innovations impedes innovation.

[22]We run this analysis at the level of SIC codes rather than NAICS codes to reduce noise because the patent-to-industry linkage is available only at the SIC level.

Whether supplier linkages are measured using the input-output network (columns 1 and 2) or the citation network (columns 3 and 4), we find a consistent negative relationship between the variance of supplier innovation and an industry's own patenting (as well as the expected strong positive relationship for the mean of supplier patenting). These relationships are typically (but not always) more precisely estimated when using the citation network (columns 3 and 4) in place of the input-output network, which appears plausible since the citation network should be more relevant for patenting, even if measured with greater imprecision as noted above.[23]

In Panels B and C, we switch from our quantity-based measure of patenting inputs in Panel A to the two quality-based measures: the citation-weighted count of patents in Panel B, and the count of top quintile patents in Panel C. Using these quality-based estimates substantially increases the precision of the key explanatory variable. In column 1 of Panel B, the variance of supplier innovation has a coefficient of -1.563 with standard error = 0.228, and in Panel C, it has a coefficient of -0.687 with standard error = 0.092 (compared to the noisier estimate of -0.750 with standard error = 0.495 in Panel A).

Appendix Table A10 demonstrates the robustness of these specifications to many of the modifications explored for the baseline results in Table 4. Specifically, similar patterns appear in weighted regressions, when we exclude the computer and electronics sector, and when we fix the upstream network, either the input-output structure or the citation network, in a given year. Results are also robust to using alternate measures of high-value patenting, such as the number of breakthrough patents identified in Kelly et al. (2018).

Overall, the results in Tables 7 and 8 bolster our interpretation that a main channel linking supplier productivity to an industry's productivity growth is innovation.

6 International Evidence

Our primary analysis focuses on productivity growth and innovation in the United States. We supplement that evidence here by estimating a variant of equation (6) for productivity growth across European countries. As outlined in Section 3, we use WIOD data to construct consistent input-output linkages for 30 industries for Spain, France, the US, Austria, Finland, the Netherlands, Italy, Germany, and the UK. We fix the global input-output table in 2000 and focus on industry TFP growth in this cross-country sample ranging from 1987 to 2007. These data enable us to include international input-output linkages, which we exploit in our

[23]In Appendix Table A9, we explore the relationship between supplier patenting and downstream TFP. Unsurprisingly given the imprecision of the crosswalk between patent classes and industry classification, these results are noisy and do not provide a clear picture.

calculations of the mean and variance of supplier TFP growth.[24]

We report these cross-country estimates in Table 10. We report our baseline specifications in the first four columns. These specifications are unweighted and all include a combination of country effects, year effects, and year-by-country effects or year-by-industry effects, as noted at the base of each column. In column 1, we focus on a specification containing country and year effects and verify that an industry's TFP growth is predicted by the average TFP growth of its suppliers. Column 2 includes the variance of supplier TFP growth. The coefficient on this measure is negative, highly significant and broadly similar to the US-based estimates: -0.824 (standard error = 0.212).

Subsequent columns probe the robustness of this finding. Column 3 includes country-by-year effects, thus exploiting only within-country variation to identify the relationship between industry TFP growth and the mean and variance of supplier TFP growth. The relationship is similar although somewhat smaller. In particular, the coefficient on the variance term is -0.535 (standard error = 0.143). Column 4 additionally includes industry-by-year interactions, so that the identifying variation is only within-industry cross-country rather than cross-industry as in the other specifications in the paper. In this demanding specification, the coefficient on the variance term remains negative and statistically significant, at -0.444 (standard error = 0.157).

The negative effect of supplier TFP variance is also present when we include the lagged dependent variable to control for mean reversion dynamics (column 5). It is weaker but still present when we use (nominal) value-added weights instead of our baseline unweighted specification (columns 6 and 7, with and without controlling for lagged dependent variable). It is equally large, and in this case statistically significant, when we focus on a 10-year panel rather than stacked five-year changes in column 8.[25] For the baseline specification, we used the full international input-output table, incorporating inputs from each country-industry pair. In column 9, we show that the estimates are similar when we only use each country's domestic input-output network.[26]

These cross-country models also enable us to investigate whether our mechanism can account for the international slowdown in productivity growth. Figure 6, which is analogous to Figure 5 for the US, reports the results of this exercise. Across the European countries in our sample, we estimate that the rising variance of supplier TFP reduced aggregate productivity growth in

[24]Specifically, we use the world input-putput tables to calculate the input share $\alpha_{ik,jl}$ as the share of inputs from industry i in country k that came from industry j in country l. The shares are based only on the nine countries listed above.

[25]See Appendix Table A11 for robustness of results for the United States aggregated to the the 30 industries used in Table 10.

[26]We do not report estimates using the manufacturing sample in this case, because manufacturing industries are not sufficiently disaggregated in this data set, and doing so would reduce our sample by about two-thirds.

eight of nine countries—all except Italy. This bottleneck effect is largest in Finland and the Netherlands, where we estimate that it reduced aggregate TFP growth over this period by 30% and 60%, respectively.

7 Conclusion

Despite the exponential pace of innovations in the information and communications technology and electronics sectors, aggregate productivity growth in the United States and many other industrialized nations has been disappointing since the 1970s, and ever more so since the early 2000s. Some have interpreted this pattern as reflecting a severe underestimation of quality and actual productivity growth or, alternatively, a lull that proceeds an inexorable surge in productivity. Others have read it as evidence of a permanent slackening of productivity growth as good ideas have become harder to find.

We propose a different hypothesis—one that implies neither a permanent slowdown in productivity growth nor an incipient surge—and investigate it empirically. The main idea driving our approach is that innovation in any one industry relies on complementary innovations in its input suppliers (or other industries linked to it). When innovation is unbalanced across industries, this holds up aggregate productivity by creating innovation bottlenecks along the input-output network.

After presenting a simple version of this productivity bottleneck hypothesis, we explore it using data on input-output linkages, citation linkages, patenting, and total factor productivity growth. Across a variety of measurement approaches, productivity outcomes, and countries, we verify the primary prediction of this hypothesis, which is that an industry's productivity growth is augmented by the mean productivity growth of its suppliers (measured by TFP or innovation) and is hampered by the variance of their productivity growth. Our primary evidence is established by using input-output linkages and TFP growth to document the sensitivity of industry productivity growth to the mean and variance of supplier productivity growth. We bolster this evidence by documenting a similar relationship when we look at supplier linkages traced via knowledge flows (the patent citation network) and when we focus on patenting behavior, rather than TFP growth, as our productivity measure.

Our estimates suggest that the bulk of the productivity slowdown in the US (and several other industrialized economies) can be accounted for by the sizable increase in the cross-industry variance of TFP growth and innovation. If, for example, the variance of TFP growth had remained at its 1977-1987 level for the subsequent two decades, we calculate that US manufacturing productivity would have grown twice as rapidly in 1997-2007 as actually occurred—yielding

a counterfactual growth rate in this decade exceeding its observed level in either of the two prior decades.

We view our paper as a first step in the theoretical and empirical investigation of the interlinked nature of innovation across sectors. Many areas of research appear fruitful based on our results. First, our hypothesis raises a critical theoretical question: will the endogenous direction of technological progress tend to clear productivity bottlenecks, or might the market mechanism exacerbate imbalances? Second, our results highlight the need for alternative empirical strategies for exploring dependencies among innovating sectors and the innovations generated by their suppliers. Third, these same relationships could be tested using firm-level data, where we suspect that the importance of supplier-customer linkages would be even larger. Fourth, it would be valuable to explore this hypothesis using historical data, focusing for example on major technological breakthroughs in the first half of the 20th century. Finally, our framework makes an important prediction whose verification awaits the passage of time: if and when lagging industries increase their innovation and productivity growth rate, a rapid takeoff in aggregate productivity can become possible.

References

Daron Acemoglu and Pablo D Azar. Endogenous production networks. *Econometrica*, 88(1): 33–82, 2020.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.

Daron Acemoglu, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012.

Daron Acemoglu, David Dorn, Gordon H Hanson, Brendan Price, et al. Return of the solow paradox? it, productivity, and employment in us manufacturing. *AER Papers and Proceedings*, 104(5):394–99, 2014.

Daron Acemoglu, Ufuk Akcigit, and William R Kerr. Innovation network. *Proceedings of the National Academy of Sciences*, 113(41):11483–11488, 2016.

Philippe Aghion and Peter Howitt. A Model of Growth through Creative Destruction. *Econometrica*, 60(2):323–351, 1992.

Philippe Aghion, Ufuk Akcigit, Antonin Bergeaud, Richard Blundell, and David Hémous. Innovation and top income inequality. *The Review of Economic Studies*, 86(1):1–45, 2019.

Ufuk Akcigit and Sina T Ates. What happened to us business dynamism? Technical report, National Bureau of Economic Research, 2019.

Dan Andrews, Chiara Criscuolo, and Peter N Gal. The best versus the rest: the global productivity slowdown, divergence across firms and the role of public policy. 2016.

Enghin Atalay. How important are sectoral shocks? *American Economic Journal: Macroeconomics*, 9(4):254–80, 2017.

David Autor, David Dorn, and Gordon H Hanson. The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–68, 2013.

David Baqaee and Emmanuel Farhi. Networks, barriers, and trade. Technical report, National Bureau of Economic Research, 2019.

Nicholas Bloom, Charles I Jones, John Van Reenen, and Michael Webb. Are ideas getting harder to find? *American Economic Review*, 110(4):1104–44, 2020.

Nick Bostrom. *Superintelligence: Paths, dangers, strategies*. Oxford University Press Cambridge, 2014.

Timothy F Bresnahan and Manuel Trajtenberg. General purpose technologies 'engines of growth'? *Journal of econometrics*, 65(1):83–108, 1995.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson. The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, forthcoming.

David M Byrne, John G Fernald, and Marshall B Reinsdorf. Does the united states have a productivity slowdown or a measurement problem? *Brookings Papers on Economic Activity*, 2016(1):109–182, 2016.

Lauren Cohen, Umit G Gurun, and Scott Duke Kominers. The growing problem of patent trolling. *Science*, 352(6285):521–522, 2016.

Tyler Cowen. *The great stagnation: How America ate all the low-hanging fruit of modern history, got sick, and will (eventually) feel better: A Penguin eSpecial from Dutton*. Penguin, 2011.

Paul A David. The dynamo and the computer: an historical perspective on the modern productivity paradox. *The American Economic Review*, 80(2):355–361, 1990.

Ryan A Decker, John Haltiwanger, Ron S Jarmin, and Javier Miranda. Declining dynamism, allocative efficiency, and the productivity slowdown. *American Economic Review*, 107(5): 322–26, 2017.

Peter H Diamandis and Steven Kotler. *Abundance: The future is better than you think*. Simon and Schuster, 2012.

Robert J Gordon. *The rise and fall of American growth: The US standard of living since the civil war*. Princeton University Press, 2017.

Gene M Grossman and Elhanan Helpman. *Innovation and growth in the global economy*. MIT press, 1991.

Jonathan Gruber and Simon Johnson. *Jump-starting America: How breakthrough science can revive economic growth and the American dream*. Hachette UK, 2019.

Elhanan Helpman and Manuel Trajtenberg. Diffusion of general purpose technologies. Technical report, National bureau of economic research, 1996.

Andrei Iancu. Remarks by director iancu at the 2019 international intellectual property conference. 2019.

Bryan Kelly, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. Measuring technological innovation over the long run. Technical report, National Bureau of Economic Research, 2018.

William R Kerr. Ethnic scientific communities and international technology diffusion. *The Review of Economics and Statistics*, 90(3):518–537, 2008.

Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2):665–712, 2017.

Ray Kurzweil. *The singularity is near: When humans transcend biology*. Penguin, 2005.

David Laws. 13 sextillion & counting: The long & winding road to the most frequently manufactured human artifact in history. 2018.

John B Jr. Long and Charles I Plosser. Real business cycles. *Journal of political Economy*, 91 (1):39–69, 1983.

Mariana Mazzucato. Financing innovation: Creative destruction vs. destructive creation. *Industrial and Corporate Change*, 22(4):851–867, 2013.

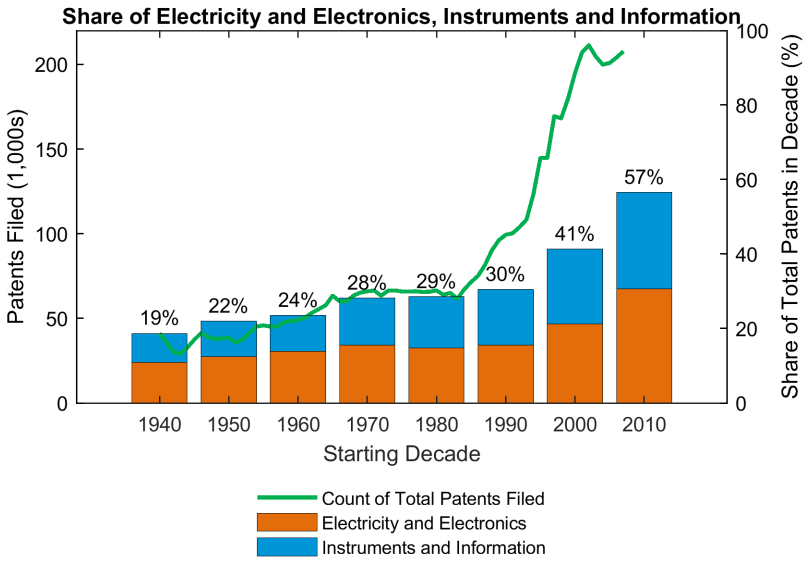
John N Park, Robert L Steigerwald, and Loren H Walker. Rectifier circuits using transistors as rectifying elements, February 24 1976. US Patent 3,940,682.

Estanislao Navarro Sosa. Electron tube. 2013.

Chad Syverson. Challenges to mismeasurement explanations for the us productivity slowdown. *Journal of Economic Perspectives*, 31(2):165–86, 2017.

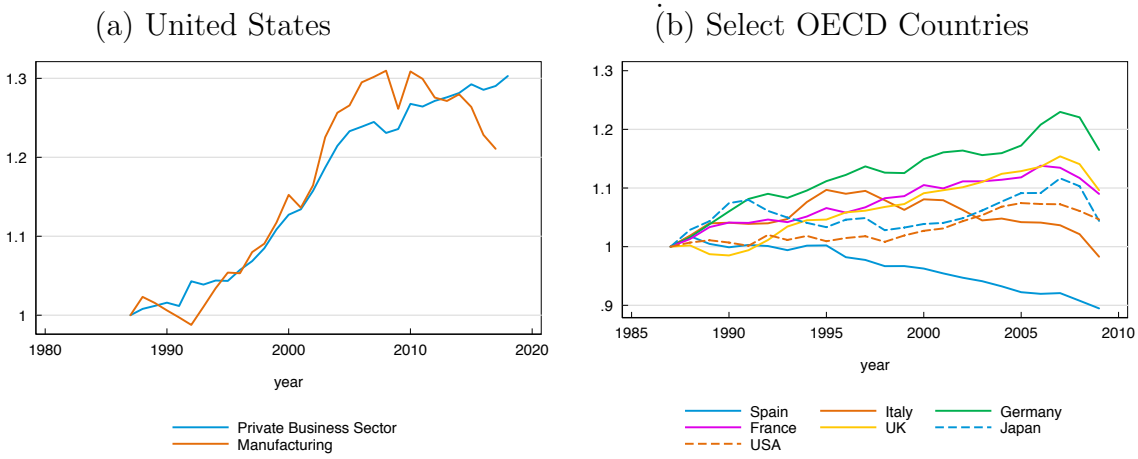
8 Tables and Figures

Figure 1: Time Series for U.S. Patents



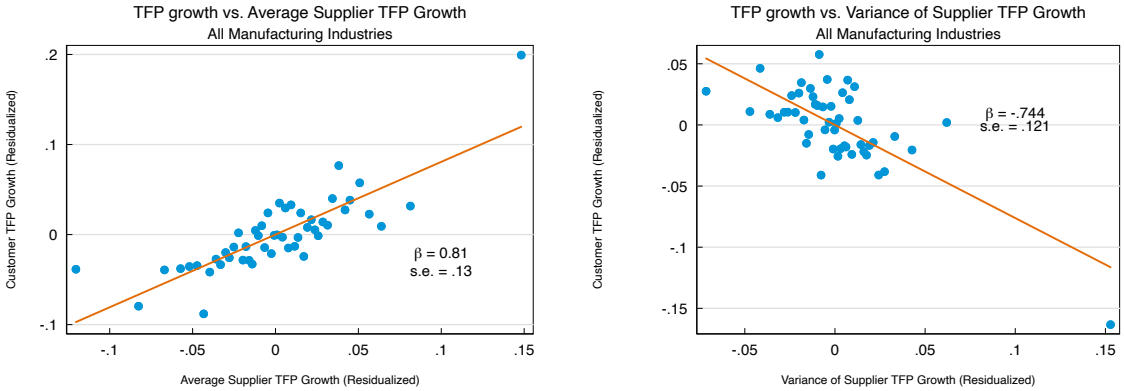
Notes: TO BE ADDED

Figure 2: Time Series for Aggregate Total Factor Productivity



Notes: Left panel shows TFP time series in the United States for 1987 - 2017. Data comes from BLS Multifactor Productivity database. Right panel shows TFP time series from 1987-2010 for the listed countries. Data for OECD countries comes from 2012 release of KLEMS.

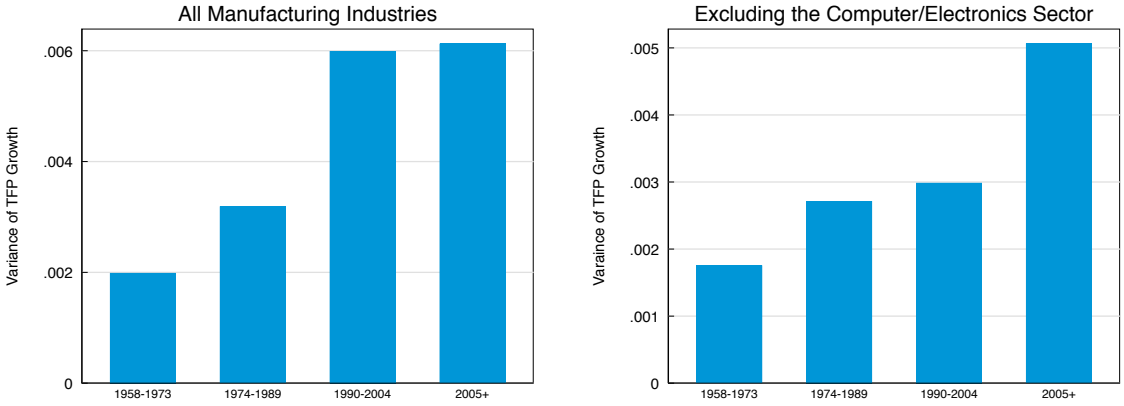
Figure 3: Bottleneck Patterns: Distribution of Upstream TFP Growth



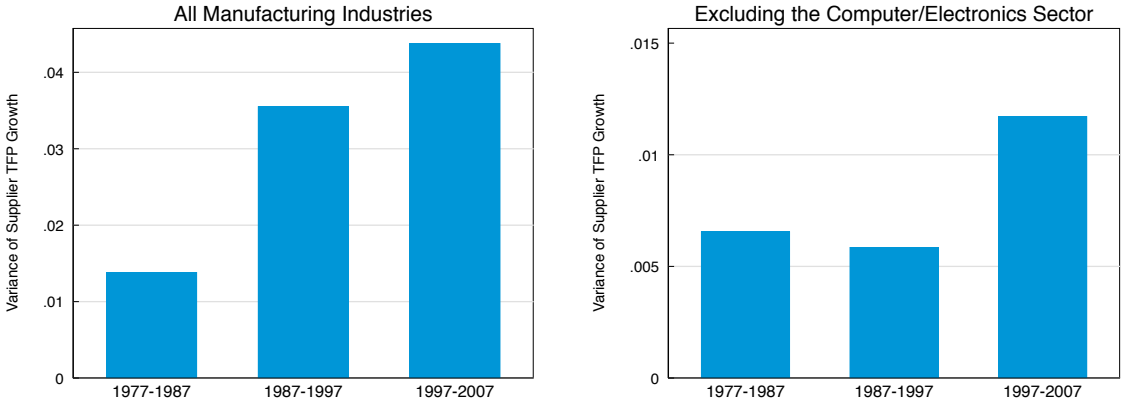
Notes: Figure reports binscatters with 50 bins for the regression model in column 2 of Panel A of Table 2 for the relationship between manufacturing TFP growth and the mean and variance of supplier TFP growth. In the left panel, variables on the x and y axis are residuals from a regression of the variable on time fixed effects and the upstream variance of TFP growth. In the right panel, variables on the x and y axis are residuals from a regression of the variable on time fixed effects and the upstream average of TFP growth. Regressions include stacked 5-year changes from 1977-2007. Industries are unweighted.

Figure 4: Variance of TFP growth

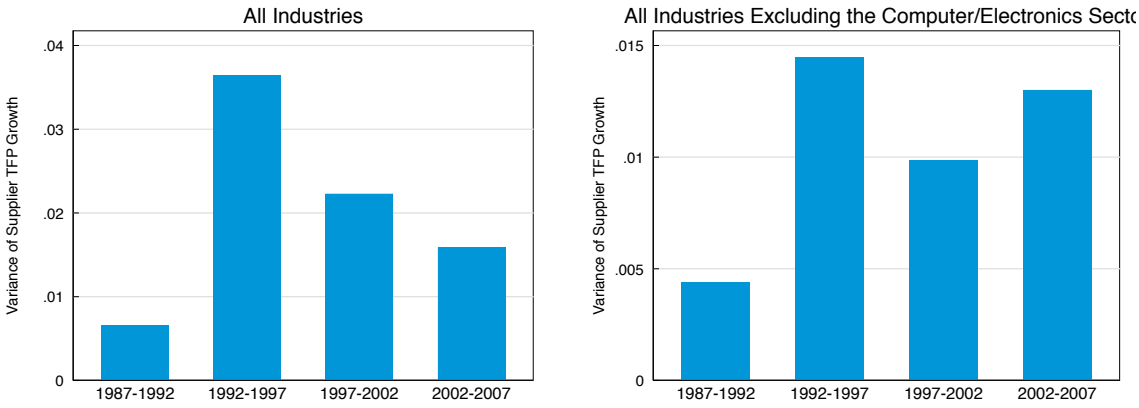
Panel A: Variance of TFP: Manufacturing Industries



Panel B: Variance of Supplier TFP: Manufacturing Industries

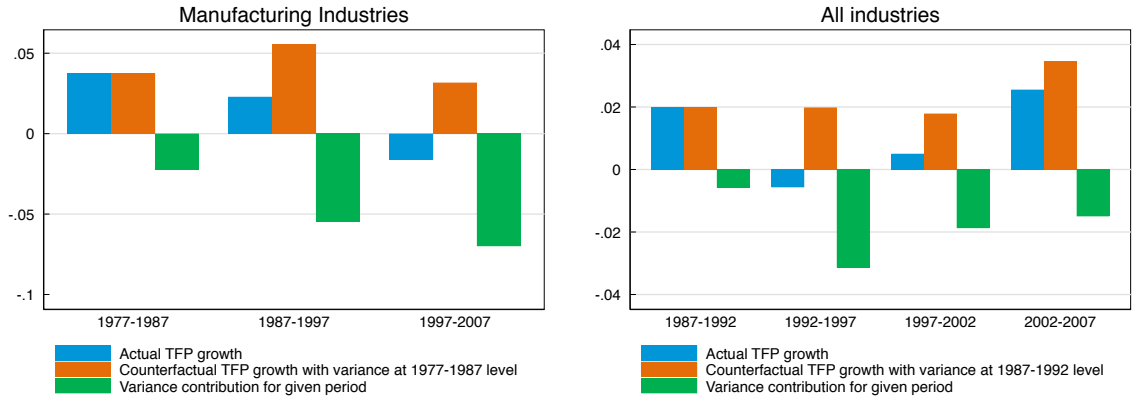


Panel C: Variance of Supplier TFP: All Industries



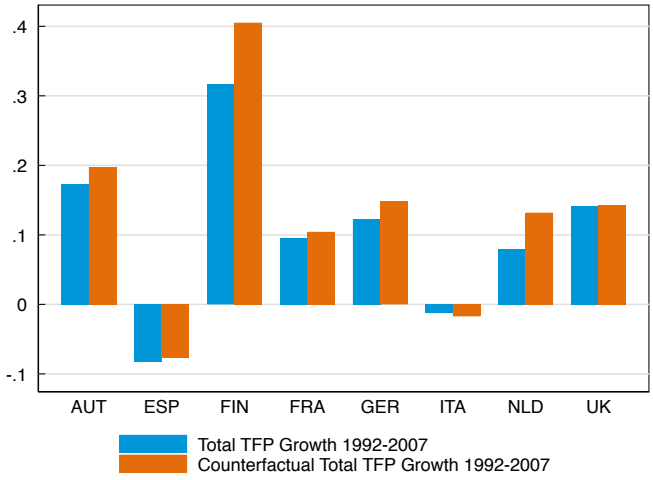
Notes: Each industry is weighted by its share of nominal value-added. For panel A, sample includes the annual change in TFP for all years from 1958-2011 and for panel B and C, sample includes stacked 5-year changes from 1977-2007. In Panel B and Panel C, the input-output network is defined at the beginning of the 5-year period. Figures on the right exclude the computer and electronics sector (NAICS code 334).

Figure 5: Magnitude of Bottleneck Estimates



Notes: Coefficients on upstream average TFP and upstream variance of TFP for the left and right panel come from Panel A and Panel B of Table 2, Column 3, respectively. Industries are aggregated using each industry's share of 1987 real value-added as the weight. In the left panel, counterfactual TFP is the TFP that would have obtained if the variance of TFP growth had been as it was from 1977-1987. This is calculated as actual TFP growth less the coefficient on the upstream variance times the difference between the variance of TFP in the given period and the variance of TFP between 1977-1987. In the right panel, the counterfactual holds constant the variance of TFP from 1987-1992.

Figure 6: Magnitude of Bottleneck Estimates in International Data



Notes: The coefficient on the upstream variance of TFP come from Table 10, Column 2. Counterfactual TFP is the TFP that would have obtained if the variance of TFP growth had been as it was from 1992-1997. This is calculated as actual TFP growth less the coefficient on the upstream variance times the difference between the variance of TFP in the given period and the variance of TFP between 1992-1997.

Table 1: Summary Statistics Table

	Downstream		Upstream Average		Upstream Variance	
	Mean	SD	Mean	SD	Mean	SD
Panel A: Manufacturing Industries						
Growth in log(Employment)	-.08	.258	-.087	.115	.027	.021
Growth in Price Index	.134	.178	.125	.172	.033	.059
Growth in log(TFP)	.018	.152	.033	.075	.022	.048
Growth in log(Patents)	.132	.19	.085	.119	.015	.012
Growth in log(Citation-weighted patents)	.139	.233	.092	.138	.02	.019
Growth in log(Number of top patents)	.27	.322	.177	.205	.04	.036
Panel B: All Industries						
Growth in log(Employment)	-.079	.266	-.084	.123	.025	.023
Growth in Price Index	.095	.147	.081	.145	.04	.069
Growth in log(TFP)	.015	.155	.034	.079	.028	.057
Panel C: International panel						
Growth in log(TFP)	.046	0.16	.041	.068	.018	.023

Notes: Panel A includes stacked 5-year changes for 462 manufacutring industries from 1977-2007. Panel B includes stacked 5-year changes for 504 industries from 1987-2007. Panel C includes stacked 5-year changes for 30 industries from 1987-2007 for 9 countries (United States, Spain, France, Austria, Finland, the Netherlands, Italy, Germany and the UK) Upstream metrics are taken using intermediate cost shares from the input-output matrix. Indiviudal industries are unweighted.

Table 2: Relationship between industry TFP growth and supplier TFP growth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				A. Manufacturing Only				
Input Average	0.425	0.810	0.653	0.676	0.255	0.837	0.617	0.193
	(0.139)	(0.130)	(0.074)	(0.170)	(0.122)	(0.133)	(0.086)	(0.165)
Input Variance		-0.744	-0.912			-0.889	-0.874	
		(0.121)	(0.118)			(0.109)	(0.101)	
Input Bottom Decile				0.059	0.378			0.396
				(0.113)	(0.091)			(0.152)
Input Top Decile				-0.110	-0.081			-0.099
				(0.033)	(0.032)			(0.039)
Ind. Fixed Effects	no	no	yes	no	yes	no	yes	yes
Industry Weighting	None	None	None	None	None	Real VA	Real VA	Real VA
Observations	2772	2772	2772	2772	2772	2772	2772	2772
R-Squared	0.108	0.133	0.371	0.118	0.361	0.168	0.434	0.425
				B. All Industries				
Input Average	0.343	0.915	0.780	0.636	0.387	0.278	0.219	-0.104
	(0.178)	(0.161)	(0.119)	(0.183)	(0.170)	(0.249)	(0.228)	(0.332)
Input Variance		-0.905	-1.087			-0.442	-0.974	
		(0.158)	(0.191)			(0.230)	(0.303)	
Input Bottom Decile				0.164	0.422			0.499
				(0.099)	(0.115)			(0.260)
Input Top Decile				-0.117	-0.139			-0.128
				(0.034)	(0.035)			(0.072)
Ind. Fixed Effects	no	no	yes	no	yes	no	yes	yes
Industry Weighting	None	None	None	None	None	Real VA	Real VA	Real VA
Observations	2016	2016	2016	2016	2016	2016	2016	2016
R-Squared	0.079	0.102	0.399	0.090	0.395	0.022	0.365	0.380

Notes: Standard errors are clustered at the industry level. In columns 7, 8 and 9, observations are weighted by the industry's 1987 share of real value-added. Regressions in Panel A include stacked 5-year changes from 1977-2007 and regressions in Panel B include stacked 5-year changes from 1987-2007. Time fixed effects are included in all specifications and industry fixed effects are included where indicated. Industries are defined using 1997 NAICS codes.

Table 3: Examples of Limiting and Limited Industries

Panel A: List of Select Fast-Growing Industries that Drive Rising TFP Variance
Semiconductor and Related Devices
Electronic Computers
Iron and Steel Mills
Computer Storage Devices
Radio and Television Broadcasting and Wireless Communications Equipment
Panel B: List of Select Bottleneck Industries
Petroleum Refineries
Pharmaceutical Preparation
Turbine and Turbine Generator Set Units
Printed Circuit Assembly
Basic Organic Chemicals
Panel C: List of Select Limited Industries
Surgical and Medical Instruments
Relay and Industrial Controls
Gasoline Engine and Engine Parts
Guided Missile and Space Vehicles
Industrial Valves

Notes: Bottleneck industries are defined as those that, were TFP in that industry to increase by 10%, the variance of TFP growth across supplying industries would fall the most. Fast growing industries are conversely defined as those that, were TFP in that industry to increase by 10%, the variance of TFP growth across supplying industries would fall the least. Limited industries are defined as those with the highest variance of TFP among suppliers among the 100 manufacturing industries with the highest value-added and further, the 50 of those top 100 with the highest average TFP growth among suppliers. Sample is restricted to 462 manufacturing industries from 1997-2007. See Appendix Table A4 for the ordered list of top 10 industries in each category over this time period.

Table 4: Robustness for Downstream TFP and Upstream TFP: Manufacturing Industries

	(1) Baseline	(2) 10-year	(3) Cov.	(4) Lagged	(5) China Shock	(6) No Comp.	(7) Fixed IO	(8) All Inputs	(9) Leads
				A: Without Industry Trends					
Input Average	0.810	0.931	0.660	0.776	0.554	0.560	0.878	1.723	
	(0.130)	(0.182)	(0.116)	(0.108)	(0.122)	(0.066)	(0.115)	(0.256)	
Input Variance	-0.744	-0.477	-0.640	-0.747	-0.718	-1.231	-0.711	-0.788	
	(0.121)	(0.094)	(0.127)	(0.115)	(0.114)	(0.587)	(0.161)	(0.441)	
Input Covariance			-0.069						
			(0.170)						
Lagged 5-year TFP				0.106					
				(0.091)					
Future Input Average									0.166
									(0.145)
Future Input Variance									0.065
									(0.098)
Observations	2772	1386	2310	2772	1386	2604	2772	2772	2772
R-Squared	0.133	0.100	0.123	0.142	0.129	0.122	0.153	0.130	0.085
				B: With Industry Trends					
Input Average	0.653	0.704	0.482	0.655	0.510	0.652	0.716	1.530	
	(0.074)	(0.122)	(0.074)	(0.078)	(0.100)	(0.072)	(0.071)	(0.173)	
Input Variance	-0.912	-0.641	-0.647	-0.908	-0.704	-1.197	-0.961	-1.562	
	(0.118)	(0.107)	(0.131)	(0.120)	(0.156)	(0.624)	(0.139)	(0.347)	
Input Covariance			-0.399						
			(0.143)						
Lagged 5-year TFP				-0.209					
				(0.043)					
Future Input Average									-0.006
									(0.079)
Future Input Variance									-0.010
									(0.101)
Observations	2772	1386	2310	2772	1386	2604	2772	2772	2772
R-Squared	0.371	0.549	0.385	0.396	0.478	0.252	0.379	0.367	0.334
				C. Weighting by Industry Real Value-Added					
Input Average	0.837	0.981	0.745	0.822	0.619	0.675	0.989	1.715	
	(0.133)	(0.242)	(0.127)	(0.110)	(0.132)	(0.092)	(0.133)	(0.279)	
Input Variance	-0.889	-0.624	-0.775	-0.898	-0.816	-1.770	-0.891	-1.225	
	(0.109)	(0.111)	(0.143)	(0.098)	(0.105)	(0.869)	(0.135)	(0.426)	
Input Covariance			-0.129						
			(0.158)						
Lagged 5-year TFP				0.161					
				(0.108)					
Future Input Average									0.120
									(0.184)
Future Input Variance									0.004
									(0.159)
Observations	2772	1386	2310	2772	1386	2604	2772	2772	2772
R-Squared	0.168	0.094	0.168	0.191	0.176	0.167	0.199	0.152	0.096

Notes: Standard errors are clustered at the industry level. All regression includes stacked 5-year changes, except for in column 2 where we include stacked 10-year changes. Time fixed effects are included in all specifications and industry fixed effects are included in Panel B. Observations are unweighted in Panel A and B and are weighted by the industry's share of 1987 real value-added in Panel C. Column 5 includes controls for the china shock in the industry as well as the input-weighted average China shock among suppliers. We construct the China shock as in Autor Dorn and Hanson (2013). Column 6 excludes the computer sector (NAICS 334) both from the regression sample and from the construction of the input mean and variance. Column 7 fixed the input-output table in 1987. Column 8 defines the input-output network to use the share not among intermediates but among all inputs. Column 9 includes the input mean and variance of TFP growth in the future 5-year period.

Table 5: Country-Specific Instruments

	(1)	(2)	(3)	(4)	(5)	(6)
			A: 2SLS Estimates			
	OLS Estimates		2SLS Estimates			
Upstream Average	1.119	0.951	1.369	1.416	1.414	1.322
	(0.312)	(0.143)	(0.363)	(0.655)	(0.648)	(0.627)
Upstream Variance	-0.907	-1.175	-0.902	-0.887	-0.882	-1.231
	(0.212)	(0.103)	(0.385)	(0.527)	(0.521)	(0.724)
Ind. Fixed Effects	no	yes	no	yes	yes	yes
Upstream Network	IO	IO	IO	IO	IO	IO
Controls for Imports					Down	Up and Down
Observations	2940	2940	2478	2478	2478	2478
R-Squared	0.136	0.373	0.105	0.350	0.351	0.369
First-Stage F-Stat	0	0	1.38	.63	.6	.51
			B: Rank of TFP growth, 2SLS Estimates			
	OLS Estimates		2SLS Estimates			
Upstream Average	1.119	0.951	0.928	1.093	1.090	1.122
	(0.312)	(0.143)	(0.338)	(0.348)	(0.334)	(0.310)
Upstream Variance	-0.907	-1.175	-0.667	-1.480	-1.483	-1.675
	(0.212)	(0.103)	(0.445)	(0.661)	(0.655)	(0.661)
Ind. Fixed Effects	no	yes	no	yes	yes	yes
Upstream Network	IO	IO	IO	IO	IO	IO
Controls for Imports					Down	Up and Down
Observations	2940	2940	2478	2478	2478	2478
R-Squared	0.136	0.373	0.117	0.374	0.374	0.375
First-Stage F-Stat	0	0	.8	2.1	2.07	1.83
			C: Rank of TFP growth, LIML Estimates			
	OLS Estimates		LIML Estimates			
Upstream Average	1.119	0.951	0.928	1.094	1.091	1.123
	(0.312)	(0.143)	(0.342)	(0.349)	(0.336)	(0.311)
Upstream Variance	-0.907	-1.175	-0.664	-1.482	-1.485	-1.677
	(0.212)	(0.103)	(0.449)	(0.665)	(0.659)	(0.665)
Ind. Fixed Effects	no	yes	no	yes	yes	yes
Upstream Network	IO	IO	IO	IO	IO	IO
Controls for Imports					Down	Up and Down
Observations	2940	2940	2478	2478	2478	2478
R-Squared	0.136	0.373	0.117	0.374	0.374	0.375
First-Stage F-Stat	0	0	.8	2.1	2.07	1.83

Notes: Standard errors are clustered at the industry level. Time fixed effects are included in all specifications. The three countries we use as instruments are France, Germany and the UK. Panels A and B report the two-stage least squares estimates and Panel C reports the limited information maximum likelihood estimates. Column 5 includes controls for the growth in an industry's imports from each country to the US and column 6 also includes the growth in imports from each country to the supplying industries.

Table 6: Exploring Neoclassical Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline		Prices		Employment		Combined	
	A: Time-Varying Input Shares							
TFP Average	0.810	0.653	0.815	0.692	0.720	0.497	0.602	0.436
	(0.130)	(0.074)	(0.113)	(0.083)	(0.134)	(0.075)	(0.107)	(0.089)
TFP Variance	-0.744	-0.912	-0.686	-0.527	-0.703	-0.786	-0.655	-0.424
	(0.121)	(0.118)	(0.232)	(0.243)	(0.118)	(0.115)	(0.233)	(0.245)
Input Price Average			0.006	0.077			-0.141	-0.069
			(0.085)	(0.061)			(0.091)	(0.065)
Input Price Variance			-0.051	-0.329			-0.123	-0.381
			(0.204)	(0.198)			(0.201)	(0.201)
Input Employment Average					0.224	0.244	0.264	0.262
					(0.045)	(0.056)	(0.051)	(0.062)
Input Employment Variance					0.166	-0.106	0.117	-0.213
					(0.219)	(0.235)	(0.218)	(0.246)
Ind. Fixed Effects	no	yes	no	yes	no	yes	no	yes
Observations	2772	2772	2772	2772	2772	2772	2772	2772
R-Squared	0.133	0.371	0.133	0.373	0.149	0.384	0.152	0.387
	B: Time-Invariant Input Shares							
TFP Average	0.878	0.716	0.840	0.741	0.802	0.574	0.669	0.512
	(0.115)	(0.071)	(0.098)	(0.078)	(0.119)	(0.071)	(0.093)	(0.082)
TFP Variance	-0.711	-0.961	-1.061	-0.760	-0.639	-0.828	-0.912	-0.535
	(0.161)	(0.139)	(0.237)	(0.225)	(0.161)	(0.131)	(0.234)	(0.217)
Input Price Average			-0.026	0.059			-0.156	-0.081
			(0.067)	(0.053)			(0.075)	(0.057)
Input Price Variance			0.314	-0.182			0.161	-0.326
			(0.239)	(0.237)			(0.232)	(0.235)
Input Employment Average					0.131	0.153	0.170	0.181
					(0.038)	(0.051)	(0.043)	(0.055)
Input Employment Variance					-0.149	-0.338	-0.150	-0.410
					(0.173)	(0.213)	(0.173)	(0.215)
Ind. Fixed Effects	no	yes	no	yes	no	yes	no	yes
Observations	2772	2772	2772	2772	2772	2772	2772	2772
R-Squared	0.153	0.379	0.154	0.379	0.163	0.389	0.167	0.391

Notes: Standard errors are clustered at the industry level. All regressions include year fixed effects and industry fixed effects are included where indicated. Sample includes stacked 5-year changes for manufacturing industries from 1977-2007. Panel A allows input shares to vary with each 5-year period. Panel B fixes the input shares at their 1987 levels. Industries are unweighted.

Table 7: Comparing Innovation Network: Input-Output and Citation Networks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	A: Time-Varying Input Shares							
Input Average (Citations)	1.421	1.143	1.137	0.853	1.397	0.947	1.030	0.702
	(0.227)	(0.274)	(0.227)	(0.294)	(0.257)	(0.311)	(0.282)	(0.310)
Input Variance (Citations)	-1.201	-1.542	-0.639	-0.909	-1.457	-2.239	-0.715	-1.458
	(0.388)	(0.706)	(0.663)	(0.831)	(0.423)	(0.679)	(0.712)	(0.945)
Input Average (IO)			0.295	0.296			0.370	0.249
			(0.113)	(0.100)			(0.130)	(0.114)
Input Variance (IO)			-0.437	-0.481			-0.515	-0.411
			(0.173)	(0.164)			(0.143)	(0.154)
Ind. Fixed Effects	no	yes	no	yes	no	yes	no	yes
Industry Weighting	None	None	None	None	Real VA	Real VA	Real VA	Real VA
Observations	1844	1844	1844	1844	1844	1844	1844	1844
R-Squared	0.125	0.408	0.133	0.415	0.159	0.474	0.174	0.481
	B: Time-Invariant Input Shares							
Input Average (Citations)	1.483	1.156	1.134	0.810	1.404	0.932	0.889	0.448
	(0.222)	(0.277)	(0.229)	(0.302)	(0.259)	(0.315)	(0.267)	(0.333)
Input Variance (Citations)	-1.225	-1.709	-0.884	-1.237	-1.371	-2.244	-0.780	-1.535
	(0.394)	(0.766)	(0.831)	(1.120)	(0.427)	(0.718)	(0.738)	(1.051)
Input Average (IO)			0.340	0.347			0.550	0.537
			(0.105)	(0.106)			(0.123)	(0.118)
Input Variance (IO)			-0.364	-0.438			-0.558	-0.575
			(0.287)	(0.243)			(0.221)	(0.205)
Ind. Fixed Effects	no	yes	no	yes	no	yes	no	yes
Industry Weighting	None	None	None	None	Real VA	Real VA	Real VA	Real VA
Observations	1844	1844	1844	1844	1844	1844	1844	1844
R-Squared	0.129	0.410	0.137	0.417	0.159	0.474	0.182	0.492
	C: Aggregated Citation Network							
Input Average (Citations)	1.656	1.381	1.407	1.120	1.637	1.154	1.286	0.912
	(0.228)	(0.284)	(0.238)	(0.310)	(0.263)	(0.346)	(0.296)	(0.341)
Input Variance (Citations)	-1.309	-1.678	-0.769	-1.117	-1.474	-2.228	-0.742	-1.476
	(0.426)	(0.808)	(0.716)	(0.931)	(0.453)	(0.781)	(0.689)	(0.951)
Input Average (IO)			0.235	0.254			0.316	0.219
			(0.116)	(0.103)			(0.126)	(0.111)
Input Variance (IO)			-0.381	-0.416			-0.470	-0.392
			(0.169)	(0.165)			(0.132)	(0.141)
Ind. Fixed Effects	no	yes	no	yes	no	yes	no	yes
Industry Weighting	None	None	None	None	Real VA	Real VA	Real VA	Real VA
Observations	1848	1848	1848	1848	1848	1848	1848	1848
R-Squared	0.137	0.416	0.143	0.421	0.172	0.480	0.185	0.486

Notes: Standard errors are clustered at the industry level. Observations in columns 5 through 8 are weighted by the industry's 1987 share of real value-added. The regressions include stacked 5-year changes from 1987-2007 and include 462 manufacturing industries. Time fixed effects are included in all specifications and industry fixed effects are included where indicated. Panel A allows input shares to vary with each 5-year period. Panel B fixes the input shares at their 1987 levels. Panel C aggregates the citation network to vary only by 4-digit industry. See Appendix Table A8 for results using just the input-output network (i.e. our baseline results) for this abbreviated time period.

Table 8: Bottleneck Patterns Using Patents

	(1)	(2)	(3)	(4)
A: Number of Patents				
Input Average	1.095	0.943	1.909	1.839
	(0.049)	(0.065)	(0.041)	(0.060)
Input Variance	-0.750	-1.155	-1.371	-1.187
	(0.495)	(0.676)	(0.293)	(0.345)
Ind. Fixed Effects	no	yes	no	yes
Upstream Network	IO	IO	Cites	Cites
Observations	1804	1804	1350	1350
R-Squared	0.438	0.717	0.740	0.851
B: Citation-Weighted Number of Patents				
Input Average	1.035	0.981	1.890	2.048
	(0.044)	(0.063)	(0.038)	(0.046)
Input Variance	-1.563	-0.804	-0.617	-0.380
	(0.228)	(0.294)	(0.191)	(0.194)
Ind. Fixed Effects	no	yes	no	yes
Upstream Network	IO	IO	Cites	Cites
Observations	1804	1804	1350	1350
R-Squared	0.713	0.836	0.882	0.931
C: Number of Top 20% Patents				
Input Average	1.190	1.163	1.879	1.805
	(0.049)	(0.062)	(0.024)	(0.036)
Input Variance	-0.687	-0.714	-0.244	-0.249
	(0.092)	(0.106)	(0.046)	(0.076)
Ind. Fixed Effects	no	yes	no	yes
Upstream Network	IO	IO	Cites	Cites
Observations	1804	1804	1350	1350
R-Squared	0.897	0.928	0.970	0.978
D: International IV Estimates				
Input Average	0.924	1.015	1.810	1.890
	(0.073)	(0.122)	(0.042)	(0.062)
Input Variance	-1.951	-2.687	-1.080	-3.288
	(0.582)	(1.219)	(0.529)	(1.591)
Ind. Fixed Effects	no	yes	no	yes
Upstream Network	IO	IO	Cites	Cites
Observations	1804	1804	1350	1350
R-Squared	0.711	0.827	0.882	0.918

Notes: The dependent variable in each regression is the growth in the patenting measure. The patenting measure indicated in the panel heading refers to both the upstream and downstream patenting variable. Panel D instruments growth in citation-weighted US patenting with growth in foreign patenting. Standard errors are clustered at the industry level. Time fixed effects are included in all specifications and industry fixed effects are included where indicated. All regressions include a control for the lagged dependent variable. Regression includes stacked 5-year changes from 1982-2002 for 490 SIC-based manufacturing industries. Top 20% patents are defined following Kogal et al (2017).

Table 9: Bottleneck Patterns Using Firm-level Patenting

	Whole Sample					Pre-2005	
Citation Average	2.407	2.876	2.093	2.463	1.811	1.220	2.673
	(0.250)	(0.267)	(0.295)	(0.265)	(0.139)	(0.077)	(0.343)
Citation Variance	-2.110	-2.166	-0.649	-2.250	-1.927	-0.988	-3.133
	(0.741)	(0.724)	(0.938)	(0.791)	(0.346)	(0.275)	(1.104)
CPC*Year FEs		✓					
Sample Restriction				✓			
Dependent Var	log	log	log	log	Avg.	Avg. w/ 0s	log
Observations	155040	154472	109274	56768	441922	1129491	87811
R-Squared	0.174	0.216	0.480	0.180	0.215	0.149	0.122
Mean Dep. Var	.307	.307	.306	.317	.308	.238	.577
Avg. Citation Variance	.059	.059	.059	.059	.058	.056	.061
Std Dev. Citation Variance	.04	.04	.039	.039	.042	.043	.035

Notes: Standard errors are clustered at the firm level. Year fixed effects included in all regressions and firm fixed effects are included where indicated

Table 10: International Patterns: Cross-Country Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Baseline				Lagged Dep. Var	VA Weight	VA Weight	10-year Changes	Within-Country IO
Upstream Average	0.258	0.270	0.108	-0.229	0.263	0.274	0.258	0.277	0.279
	(0.075)	(0.080)	(0.080)	(0.113)	(0.085)	(0.106)	(0.084)	(0.116)	(0.082)
Upstream Variance		-0.824	-0.535	-0.444	-0.822	-0.563	-0.751	-0.725	-0.716
		(0.212)	(0.143)	(0.157)	(0.182)	(0.585)	(0.422)	(0.363)	(0.215)
Year FEs	X	X			X	X	X	X	X
Country FEs	X	X	X		X	X	X	X	X
Year*Country FEs				X					
Year*Industry FEs			X	X					
Lagged Dep. Var.					X		X		
Observations	982	982	982	982	896	982	896	462	982
R-Squared	0.065	0.076	0.363	0.401	0.118	0.063	0.197	0.120	0.075

Notes: Standard errors are clustered at the industry level. All regressions (except for column 8) stacked 5-year changes from 1987-2007 for 30 industries and 9 countries: Spain, France, the US, Austria, Finland, the Netherlands, Italy, Germany and the UK. Column 5 includes a control for the lagged change in TFP growth of the downstream industry (i.e. the lagged dependent variable. Column 6 weights each observation by their share of within-country value-added. In this specification, each country gets the same weight but industries are weighted by their size within each country. Column 7 both weights industries and includes a control for the lagged dependent variable. Column 8 uses 10-year stacked changes instead of 5-year stacked changes. In all regressions, the input-output table is defined using 2000 world input-output tables. In columns 1-8, the input share is defined across industries and countries. In column 9, the input share is defined only across industries (averaging across all countries).

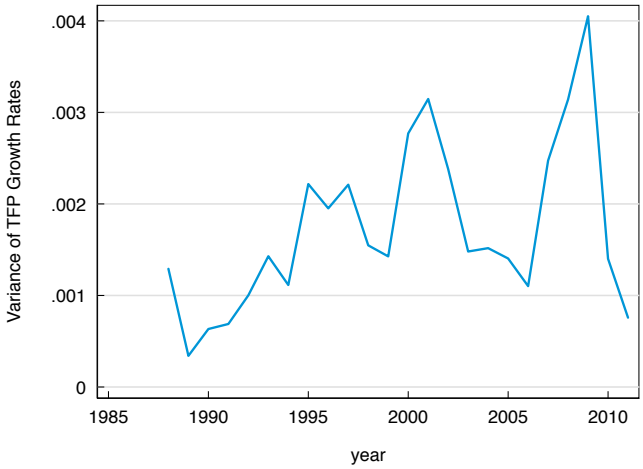
9 Data Appendix

9.1 Constructing Time-Consistent Industry Codes

In order to facilitate comparisons over time, we convert each input-output table to 1997 NAICS codes that correspond to 6-digit NAICS codes within manufacturing and roughly 3-digit NAICS codes outside of manufacturing (i.e. the level of aggregation available in the BLS multiproductivity database). The Bureau of Economic Analysis releases the detailed Input-Output tables in year-specific codes. In all years, the data is more detailed within manufacturing than in non-manufacturing sectors.The Bureau of Economic Analysis provides a mapping between these aggregate codes and that year's SIC or NAICS codes (i.e. the mapping between the 1987 IO table codes and the 1987 SIC codes). In order to preserve the highest level of disaggregation within manufacturing, when industry detail in the IO table is less than either 4-digit SIC codes or 6-digit NAICS codes, we apportion each industry's commodity purchases to the associated SIC codes on the basis of the sub-industry's material/energy costs (which we observe for all manufacturing industries in the NBER-ASM data). Similarly, we apportion each industry's production based on the sub-industry's share of shipments. Finally, using weighted industry mappings between SIC and NAICS codes over time, we convert each input-output table to 1997 NAICS codes.

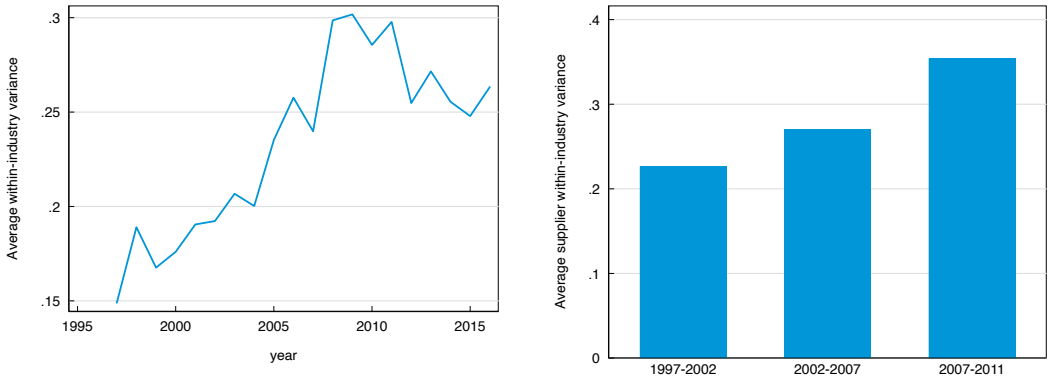
10 Appendix Tables and Figures

Figure A1: Variance of TFP Growth Across All Industries



Notes: Each industry is weighted by its share of nominal value-added. Sample includes all years from 1987-2011. There are 70 industries, corresponding to the detailed manufacutring and non-manufacturing industries in the BLS Multifactor Productivity data.

Figure A2: Variance of TFP Within Industries



Notes: Within-industry variance comes from Dispersion Statistics on Productivity (DiSP) provided by the U.S. Census Bureau. Left panel shows the average within-industry cross-firm variance in TFP, where the average is taken over 4-digit NAICS industries. Right panel shows the average within-industry variance among supplying industries (i.e. weighting by the input share of the industry). The average is taken over each 5-year period, weighting using industries by their share of value-added.

Table A1: Robustness of Bottleneck Patterns to Outliers

	(1)	(2)	(3)	(4)	(5)	(6)
			B: Manufacturing Industries			
Input Average	0.573	0.600	0.592	0.597	0.831	0.672
	(0.061)	(0.060)	(0.043)	(0.048)	(0.130)	(0.074)
Input Variance	-1.107	-2.515	-0.774	-0.757	-0.757	-0.902
	(0.719)	(0.747)	(0.059)	(0.068)	(0.120)	(0.117)
Input Skewness					0.005	0.004
					(0.002)	(0.002)
Ind. Fixed Effects	no	yes	no	yes	no	yes
Industry Weighting	None	None	None	None	None	None
Estimator	OLS	OLS	Robust	Robust	OLS	OLS
Observations	2772	2772	2772	2772	2772	2772
R-Squared	0.163	0.317	0.208	0.487	0.135	0.371
			B: All Industries			
Input Average	0.700	0.691	0.776	0.696	0.946	0.813
	(0.087)	(0.095)	(0.075)	(0.085)	(0.159)	(0.123)
Input Variance	-0.817	-2.750	-0.984	-1.139	-0.957	-1.118
	(0.927)	(1.184)	(0.087)	(0.104)	(0.166)	(0.197)
Input Skewness					0.004	0.004
					(0.002)	(0.003)
Ind. Fixed Effects	no	yes	no	yes	no	yes
Industry Weighting	None	None	None	None	None	None
Estimator	OLS	OLS	Robust	Robust	OLS	OLS
Observations	2016	2016	2016	2016	2016	2016
R-Squared	0.135	0.346	0.170	0.589	0.104	0.399

Notes: Columns 1, 2, 7 and 8 winsorizes industry TFP growth at the top and bottom 5 percent. The upstream mean and variance are re-calculated on these winsorized values. Coluns 3, 4, 9 and 10 use un-winsorized TFP codes and implement outlier-robust regressions (rreg command in STATA). Columns 5, 6, 11, and 12 include controls for the skewness of upstream TFP growth. All regression includes stacked 5-year changes from 1977-2007 in manufacturing (columns 1-4) and from 1987-2007 for all industries (columns 5-8). Time fixed effects are included in all specifications and industry fixed effects are included where specified. Standard errors are clustered at the industry level.

Table A2: Robustness of Bottleneck Patterns: Value-Added per Worker and Lagged TFP

	Manufacturing				All Industries			
	Panel A: Value-Added Per Worker							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Input Average	0.565	0.439	0.540	0.374	0.355	0.134	0.152	0.297
	(0.092)	(0.062)	(0.104)	(0.035)	(0.150)	(0.126)	(0.144)	(0.066)
Input Variance	0.009	-0.290	-0.391	-0.280	0.202	-0.154	-0.557	-0.360
	(0.098)	(0.085)	(0.127)	(0.060)	(0.125)	(0.133)	(0.301)	(0.083)
Ind. Fixed Effects	no	yes	yes	yes	no	yes	yes	yes
Industry Weight	None	None	Real VA	None	None	None	Real VA	None
Estimator	OLS	OLS	OLS	Robust	OLS	OLS	OLS	Robust
Observations	2772	2772	2772	2772	2016	2016	2016	2016
R-Squared	0.194	0.384	0.439	0.586	0.044	0.328	0.321	0.566
	Panel A: Lagged TFP							
Input Average	0.118	-0.122	-0.139	-0.214	0.048	-0.207	0.081	-0.225
	(0.134)	(0.074)	(0.076)	(0.081)	(0.125)	(0.094)	(0.296)	(0.106)
Input Variance	-0.069	-0.172	-0.085	0.214	-0.102	-0.300	-0.096	0.041
	(0.115)	(0.110)	(0.126)	(0.109)	(0.140)	(0.145)	(0.262)	(0.145)
Ind. Fixed Effects	no	yes	yes	yes	no	yes	yes	yes
Industry Weight	None	None	Real VA	None	None	None	Real VA	None
Estimator	OLS	OLS	OLS	Robust	OLS	OLS	OLS	Robust
Observations	2310	2310	2310	2310	2016	2016	2016	2016
R-Squared	0.086	0.360	0.425	0.551	0.069	0.380	0.343	0.572

Notes: All regressions include year fixed effects and stacked 5-year changes. Columns 1 through 4 include all manufacturing industries from 1977-2007 and columns 4 through 8 include all industries from 1987-2007. Columns 4 and 8 show estimates using outlier-robust regressions (using rreg command in STATA). In panel A, the independent variables are the input-weighted mean and variance of the log of growth in real value-added per worker among suppliers and the dependent variable is an industry's growth in log of real value-added per worker. In panel B, the independent variables are the input-weighted mean and variance of TFP growth among suppliers in the prediouc 5-year period and the dependent variable is an industry's growth in log of real value-added per worker in the current period.

Table A3: Robustness of Bottleneck Patterns to Including Within-Industry Variance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Upstream Average	0.509	0.090	0.202	-0.116	0.568	0.038	0.168	-0.290
	(0.166)	(0.179)	(0.169)	(0.206)	(0.174)	(0.184)	(0.168)	(0.259)
Upstream Variance	-1.132		-1.502		-1.142		-1.379	
	(0.192)		(0.247)		(0.208)		(0.318)	
Upstream Within Industry Variance	0.265	0.031	0.010	-0.351	0.123	-0.110	-0.569	-0.686
	(0.126)	(0.119)	(0.285)	(0.294)	(0.147)	(0.133)	(0.356)	(0.356)
Ind. Fixed Effects	no	no	yes	yes	no	no	yes	yes
Industry Weighting	None	None	None	None	Real VA	Real VA	Real VA	Real VA
Observations	924	924	924	924	924	924	924	924
R-Squared	0.135	0.092	0.597	0.553	0.191	0.118	0.636	0.562

Notes: Notes: Within-industry variance comes from Dispersion Statistics on Productivity (DiSP) provided by the U.S. Census Bureau. Sample includes 5-year changes from 1997-2002 and 2002-2007. Within-industry metrics are provided at the 4-digit NAICS level and are averaged across all years within the 5-year window. All regressions include year fixed effects.

Table A4: Top 10 Limiting and Limited Industries

Panel A: List of Fastest-Growing Industries

1997-2002 Industries	*2002-2007 Industries*
Semiconductor and Related Devices	Semiconductor and Related Devices
Electronic Computers	Electronic Computers
Paper (except Newsprint) Mills	Computer Storage Devices
Other Animal Food	Biological Product
Soybean Processing	Sawmills
Motor Vehicle Electrical and Electronic Equip.	In-Vitro Diagnostic Substances
Poultry Processing	All Other Basic Inorganic Chemical Manufacturing
Iron and Steel Mills	Petrochemicals
Flavoring Syrup and Concentrate	Radio/TV Broadcasting & Wireless Comm. Equip.

Panel B: List of Bottleneck Industries

1997-2002 Industries	*2002-2007 Industries*
Commercial Lithographic Printing	Petroleum Refineries
All Other Basic Organic Chemical	Pharmaceutical Preparation
Printed Circuit Assembly (Electronic Assembly)	Other Communication and Energy Wires
Corrugated and Solid Fiber Boxes	Manifold Business Forms Printing
Petrochemicals	Corrugated and Solid Fiber Boxes
Radio/TV Broadcasting & Wireless Comm. Equip.	Printed Circuit Assembly (Electronic Assembly)
Bare Printed Circuit Boards	Turbine and Turbine Generator Set Units
Other Electronic Components	Motor Vehicle Electrical and Electronic Equip.
Petroleum Refineries	Medicinal and Botanical Manufacturing
Paint and Coatings	Unsupported Plastics Film and Sheets

Panel C: List of Limited Industries

1997-2002 Industries	*2002-2007 Industries*
Photographic and Photocopying Equipment	In-Vitro Diagnostic Substances
Relay and Industrial Control	Medicinal and Botanical
Sawmills	Guided Missile and Space Vehicles
Surgical and Medical Instruments	Wineries
Guided Missile and Space Vehicles	Petroleum Refineries
All Other Motor Vehicle Parts Manufacturing	All Other Basic Organic Chemicals
Motor Vehicle Transmission and Power Train Parts	Other Commercial and Service Industry Machinery
Gasoline Engine and Engine Parts	Cement
Motor Vehicle Metal Stamping	Relay and Industrial Controls
Motor Vehicle Electrical and Electronic Equip.	Industrial Valves

Notes: Bottleneck industries are defined as those that, were TFP in that industry to increase by 10%, the variance of TFP growth across supplying industries would fall the most. Fast growing industries are conversely defined as those that, were TFP in that industry to increase by 10%, the variance of TFP growth across supplying industries would fall the least. Limited industries are defined as those with the highest variance of TFP among suppliers among the 100 manufacturing industries with the highest value-added and further, the 50 of those top 100 with the highest average TFP growth among suppliers. Sample is restricted to 462 manufacturing industries.

Table A5: Robustness for Downstream TFP and Upstream TFP: All Industries

	(1) Baseline	(2) 10-year	(3) Cov.	(4) Lagged	(5) China Shock	(6) No Comp.	(7) Fixed IO	(8) All Inputs	(9) Leads
				A: Without Industry Trends					
Input Average	0.915	1.093	0.909	0.915	0.934	0.699	0.996	1.278	
	(0.161)	(0.233)	(0.163)	(0.151)	(0.168)	(0.108)	(0.126)	(0.248)	
Input Variance	-0.905	-0.638	-0.798	-0.923	-1.026	-2.029	-0.845	-1.014	
	(0.158)	(0.117)	(0.176)	(0.152)	(0.159)	(0.710)	(0.207)	(0.382)	
Input Covariance			-0.228						
			(0.204)						
Lagged 5-year TFP				0.070					
				(0.115)					
Future Input Average									-0.003
									(0.178)
Future Input Variance									0.244
									(0.133)
Observations	2016	1008	2016	1974	1512	1904	2016	2016	2016
R-Squared	0.102	0.070	0.103	0.107	0.116	0.103	0.119	0.090	0.073
				B: With Industry Trends					
Input Average	0.780	0.892	0.687	0.724	0.896	0.770	0.876	1.169	
	(0.119)	(0.173)	(0.118)	(0.117)	(0.134)	(0.122)	(0.124)	(0.194)	
Input Variance	-1.087	-1.024	-0.907	-0.889	-1.200	-1.490	-1.139	-1.565	
	(0.191)	(0.251)	(0.196)	(0.173)	(0.214)	(0.767)	(0.204)	(0.402)	
Input Covariance			-0.760						
			(0.196)						
Lagged 5-year TFP				-0.362					
				(0.045)					
Future Input Average									-0.171
									(0.110)
Future Input Variance									0.343
									(0.158)
Observations	2016	1008	2016	1974	1512	1904	2016	2016	2016
R-Squared	0.399	0.665	0.406	0.464	0.471	0.291	0.409	0.392	0.375
				C. Weighting by Industry Real Value-Added					
Input Average	0.278	0.466	0.276	0.304	0.289	0.167	0.498	0.502	
	(0.249)	(0.376)	(0.253)	(0.255)	(0.281)	(0.267)	(0.231)	(0.454)	
Input Variance	-0.442	-0.264	-0.508	-0.474	-0.480	-3.076	-0.735	-0.947	
	(0.230)	(0.209)	(0.311)	(0.227)	(0.261)	(1.232)	(0.234)	(0.405)	
Input Covariance			0.162						
			(0.488)						
Lagged 5-year TFP				0.050					
				(0.123)					
Future Input Average									0.352
									(0.229)
Future Input Variance									0.077
									(0.242)
Observations	2016	1008	2016	1974	1512	1904	2016	2016	2016
R-Squared	0.022	0.015	0.022	0.026	0.025	0.057	0.035	0.023	0.036

Notes: Standard errors are clustered at the industry level. All regression includes stacked 5-year changes, except for in column 2 where we include stacked 10-year changes. Time fixed effects are included in all specifications and industry fixed effects are included in Panel B. Observations are unweighted in Panel A and B and are weighted by the industry's share of 1987 real value-added in Panel C. Column 5 includes controls for the china shock in the industry as well as the input-weighted average China shock among suppliers. We construct the China shock as in Autor Dorn and Hanson (2013). Column 6 excludes the computer sector (NAICS 334) both from the regression sample and from the construction of the input mean and variance. Column 7 fixed the input-output table in 1987. Column 8 defines the input-output network to use the share not among intermediates but among all inputs. Column 9 includes the input mean and variance of TFP growth in the future 5-year period.

Table A6: First Stage for Country-Specific Instruments

	Partial First Stage		First Stage		Partial First Stage		First Stage	
Dependent Variable:	Average	Variance	Average	Variance	Average	Variance	Average	Variance
			A: Level of TFP growth					
Upstream Average France	0.106		0.084	0.208	0.015		-0.020	0.095
	(0.172)		(0.168)	(0.144)	(0.114)		(0.101)	(0.073)
Upstream Average Germany	0.064		0.121	-0.039	0.052		0.131	-0.054
	(0.021)		(0.031)	(0.021)	(0.047)		(0.049)	(0.048)
Upstream Average UK	0.117		0.127	0.002	0.097		0.113	-0.053
	(0.066)		(0.061)	(0.022)	(0.077)		(0.076)	(0.049)
Upstream Variance France		0.395	0.312	0.228		0.385	0.429	0.114
		(0.287)	(0.362)	(0.210)		(0.232)	(0.338)	(0.088)
Upstream Variance Germany		-0.029	0.075	-0.040		-0.002	0.125	-0.045
		(0.021)	(0.038)	(0.024)		(0.005)	(0.044)	(0.038)
Upstream Variance UK		-0.477	-0.067	-0.027		-0.740	-0.260	-0.481
		(0.378)	(0.521)	(0.147)		(0.422)	(0.690)	(0.284)
Ind. Fixed Effects	no	no	no	no	yes	yes	yes	yes
Observations	2520	2520	2520	2520	2520	2520	2520	2520
R-Squared	0.250	0.091	0.256	0.144	0.524	0.524	0.533	0.534
			B: Rank of TFP growth					
Upstream Average France	-0.052		-0.042	-0.323	0.109		0.116	-0.175
	(0.265)		(0.254)	(0.219)	(0.166)		(0.170)	(0.120)
Upstream Average Germany	-0.344		-0.392	0.057	-0.435		-0.454	0.061
	(0.087)		(0.073)	(0.040)	(0.106)		(0.097)	(0.063)
Upstream Average UK	-0.120		-0.115	-0.089	-0.096		-0.079	0.008
	(0.109)		(0.108)	(0.057)	(0.111)		(0.114)	(0.053)
Upstream Variance France		0.032	0.038	0.037		0.040	0.031	0.037
		(0.022)	(0.032)	(0.026)		(0.028)	(0.033)	(0.025)
Upstream Variance Germany		-0.010	-0.024	-0.012		-0.024	-0.024	-0.027
		(0.012)	(0.020)	(0.012)		(0.027)	(0.041)	(0.027)
Upstream Variance UK		-0.017	-0.022	-0.012		-0.015	-0.018	-0.015
		(0.016)	(0.021)	(0.013)		(0.014)	(0.022)	(0.014)
Ind. Fixed Effects	no	no	no	no	yes	yes	yes	yes
Observations	2520	2520	2520	2520	2520	2520	2520	2520
R-Squared	0.265	0.106	0.285	0.171	0.551	0.539	0.563	0.550

Notes: Panel A specifies the upstream average and variance using TFP growth. In panel B, we rank industries in each country according to their TFP growth and calculate the input-share weighted mean and variance of TFP ranks. In panel B, the dependent variables are multiplied by 100. Standard errors are clustered at the industry level. Time fixed effects are included in all specifications and industry fixed effects are included in columns 5-8.

Table A7: Robustness of Country-specific Instruments

	Ranked Baseline		Fixed IO		Dropping UK		Dropping Germany	
Upstream Average	0.928	1.093	1.148	1.344	0.972	1.138	0.415	0.917
	(0.338)	(0.348)	(0.286)	(0.245)	(0.352)	(0.393)	(0.803)	(0.874)
Upstream Variance	-0.667	-1.480	-0.281	-0.918	-0.658	-1.480	-0.225	-1.331
	(0.445)	(0.661)	(0.519)	(0.573)	(0.426)	(0.683)	(0.817)	(0.852)
Ind. Fixed Effects	no	yes	no	yes	no	yes	no	yes
Upstream Network	IO	IO	IO	IO	IO	IO	IO	IO
Controls for Imports								
Observations	2478	2478	2478	2478	2478	2478	2478	2478
R-Squared	0.117	0.374	0.123	0.375	0.116	0.373	0.103	0.377
First-Stage F-Stat	.8	2.1	3.97	3.25	1.02	2.72	1.93	.54

Notes: Standard errors are clustered at the industry level. All regressions are estimated using 2SLS. Columns 1 through 4 and columns 7 and 8 use TFP growth in France, Germany and the UK as instruments while columns 5 and 6 use TFP growth in France and Germany. Time fixed effects are included in all specifications and industry fixed effects are included where specified.

Table A8: Comparing Innovation Network: Input-Output Specification on Modified Sample

	(1)	(2)	(3)	(4)
A: Time-Varying Input Shares				
Input Average (IO)	0.197	0.580	0.123	0.465
	(0.130)	(0.121)	(0.088)	(0.088)
Input Variance (IO)		-0.635		-0.763
		(0.114)		(0.134)
Ind. Fixed Effects	no	no	yes	yes
Industry Weighting	None	None	None	None
Observations	1848	1848	1848	1848
R-Squared	0.087	0.108	0.384	0.405
B: Time-Invariant Input Shares				
Input Average (IO)	0.343	0.640	0.339	0.521
	(0.130)	(0.107)	(0.102)	(0.099)
Input Variance (IO)		-0.608		-0.813
		(0.153)		(0.152)
Ind. Fixed Effects	no	no	yes	yes
Industry Weighting	None	None	None	None
Observations	1848	1848	1848	1848
R-Squared	0.098	0.115	0.390	0.407

Notes: Standard errors are clustered at the industry level. The regressions include stacked 5-year changes from 1987-2007 and include 462 manufacturing industries. Time fixed effects are included in all specifications and industry fixed effects are included where indicated.

Table A9: Bottleneck Patterns Using Patents

	(1)	(2)	(3)	(4)
A: Number of Patents				
Input Average	0.007	-0.034	-0.008	-0.076
	(0.004)	(0.021)	(0.027)	(0.188)
Input Variance	0.003	0.018	-0.011	-0.189
	(0.002)	(0.009)	(0.012)	(0.089)
Ind. Fixed Effects	no	yes	no	yes
Upstream Network	IO	IO	Cites	Cites
Observations	1804	1804	900	900
R-Squared	0.113	0.468	0.110	0.762
B: Citation-Weighted Number of Patents				
Input Average	0.006	-0.041	-0.007	-0.154
	(0.004)	(0.019)	(0.025)	(0.146)
Input Variance	0.002	0.016	-0.008	-0.095
	(0.002)	(0.009)	(0.011)	(0.066)
Ind. Fixed Effects	no	yes	no	yes
Upstream Network	IO	IO	Cites	Cites
Observations	1804	1804	900	900
R-Squared	0.112	0.470	0.109	0.760
C: Number of Top 20% Patents				
Input Average	0.006	-0.035	-0.014	-0.079
	(0.003)	(0.016)	(0.020)	(0.064)
Input Variance	0.003	0.012	-0.013	-0.014
	(0.001)	(0.007)	(0.009)	(0.035)
Ind. Fixed Effects	no	yes	no	yes
Upstream Network	IO	IO	Cites	Cites
Observations	1804	1804	900	900
R-Squared	0.113	0.470	0.111	0.759

Notes: Dependent variable in each regression is the TFP growth in the industry. Standard errors are clustered at the industry level. Time fixed effects are included in all specifications and industry fixed effects are included where indicated. All regressions include a control for the lagged dependent variable. Regression includes stacked 5-year changes from 1982-2002 for 490 SIC-based manufacturing industries.

Table A10: Robustness of Bottleneck Patterns Using Patents

	(1)	(2)	(3)	(4)
	A: Alternate Specifictions			
	Value-Added Wgt.	No . Comp.	Fixed Network	Fixed Network
Input Average	1.133	0.974	0.947	2.061
	(0.108)	(0.068)	(0.073)	(0.052)
Input Variance	-0.915	-1.079	-0.958	-0.345
	(0.442)	(0.291)	(0.331)	(0.227)
Ind. Fixed Effects	yes	yes	yes	yes
Upstream Network	IO	IO	IO	Cites
Observations	1804	1588	1804	1800
R-Squared	0.858	0.837	0.830	0.929
	B: Alternate Patenting Variables			
	Total Citations	Alt. Mapping	Breakthrough Patents	Foreign Patents
Input Average	1.895	1.599	1.925	2.119
	(0.038)	(0.167)	(0.038)	(0.050)
Input Variance	-0.592	0.286	-0.467	-2.912
	(0.196)	(0.448)	(0.090)	(0.710)
Ind. Fixed Effects	no	no	no	no
Upstream Network	Cites	Cites	Cites	Cites
Observations	1350	1350	1350	1350
R-Squared	0.883	0.463	0.885	0.904
	C: Patent Class Codes			
	Baseline	CPC Trends	Extended Sample	Regression Weights
Input Average	1.301	1.378	1.487	1.347
	(0.074)	(0.110)	(0.097)	(0.111)
Input Variance	-0.990	-0.830	-0.956	-0.872
	(0.283)	(0.360)	(0.309)	(0.317)
CPC Fixed Effects	no	yes	yes	yes
Upstream Network	Cites	Cites	Cites	Cites
Observations	3099	3096	4323	3099
R-Squared	0.294	0.465	0.403	0.780

Notes: Dependent variable in each regression is the growth in the patenting measure. Standard errors are clustered at the industry level. Time fixed effects are included in all specifications and industry fixed effects are included where indicated. All regressions include a control for the lagged dependent variable. Regression includes stacked 5-year changes from 1982-2002 for 490 SIC-based manufacturing industries. The patenting measure in Panel A is citation-weighted patents. Panel C includes 980 4-digit CPC codes (i.e. patent technology classes). We combine all patents applied for in roughly 5 year-intervals (i.e. 1976-1980, 1980-1985, 1985-1990, 1990-1995, 1995-2000, 2000-2005, 2005-2010, 2010-2016) and the dependent variable is the change in log patents across subsequent 5-year intervals. Columns 1, 2 and 4 in panel C include data from 1976-2005 while column 3 in panel C includes data from 1976-2016. In column 4 in panel C, each observation is weighted by the CPC codes share of total patenting in the sample.

Table A11: Aggregates Results for United States

Upstream Average	0.829	0.969	0.929	0.972	0.734	0.670	0.731	0.758
	(0.555)	(0.559)	(0.671)	(0.592)	(0.411)	(0.326)	(0.453)	(0.405)
Upstream Variance	-5.386	-5.735	-4.966	-5.503	-5.939	-5.243	-6.042	-4.597
	(2.490)	(2.249)	(3.548)	(3.312)	(1.491)	(0.927)	(2.356)	(1.352)
Year FEs	X	X	X	X	X	X	X	X
Industry FEs					X	X	X	X
Industry Weight	None	None	Value-Added	Value-Added	None	None	Value-Added	Value-Added
Lagged Dep. Var.?	No	Yes	No	Yes	No	Yes	No	Yes
Observations	116	87	116	87	116	87	116	87
R-Squared	0.112	0.186	0.105	0.149	0.542	0.696	0.512	0.671

Notes: Standard errors are clustered at the industry level. All regressions include stacked 5-year changes from 1987-2007 for 30 industries within the US.