Bottlenecks: Sectoral Imbalances and the US Productivity Slowdown*

Daron Acemoglu — MIT and NBER
David Autor — MIT and NBER
Christina Patterson — Chicago Booth and NBER

April 5, 2023

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 slow growth stems in part from an 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 that is concentrated in a subset of sectors may create bottlenecks and consequently fail to translate into commensurate aggregate productivity gains. Using data on input-output linkages, citation linkages, industry productivity growth and patenting, we find evidence consistent with this hypothesis: the variance of suppliers' Total Factor Productivity 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 United States (and several other industrialized economies) can be accounted for by a sizable increase in cross-industry variance of TFP growth and innovation. For example, if TFP growth variance had remained at the 1977–1987 level, US manufacturing productivity would have grown twice as rapidly in 1997–2007 as it did—yielding a counterfactual growth rate that would have exceeded that of 1977–1987 and 1987–1997.

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

*We thank Nick Bloom, Steve Davis, Hugo Hopenhayn, Bill Kerr, Jennifer L'aO, Valerie Ramey, and John Van Reenen for helpful comments, Rebecca Jackson, Zhe Fredric Kong, and Austin Lentsch for expert research assistance, and Ernest Liu and Song Ma for sharing international patenting data. Acemoglu and Autor acknowledge support from the Hewlett Foundation, Google, and the Smith Richardson Foundation. Acemoglu also acknowledges support from the National Science Foundation, the Sloan Foundation, and the Toulouse Network on Information Technology. Autor also acknowledges support from 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, despite 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, a rapid takeoff in the total number of patents in the 1980s; and second, a surge in the share of ICT and electronics patents during the same time interval. Between 1990 and 2010, the total number of patents granted rose from 99,000 to 208,000, while the combined number of ICT and electronics patents granted increased by approximately 87,000, accounting for the bulk of the increase. Figure 2 depicts the growth rate of Total Factor Productivity (TFP) in the US economy and in the leading OECD economies in recent decades. Productivity growth in the United States has been minimal since the mid-2000s, and it has been slower still in many OECD countries, with the possible exception of Germany.

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

This paper offers a potential reconciliation of these 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 back aggregate productivity. We propose a simple framework in which the development of new technologies or products in a given sector requires simultaneous improvements in the quality of several inputs. For example, breakthroughs in automotive technology cannot be achieved solely with improvements in engine management

[1] The orange bars correspond to the share of US Patent and Trademark Office (USPTO) patents granted in Electricity and Electronics (i.e., "electronics"), while the blue bars plot the share of patents granted in Instruments and Information (i.e., "ICT"). The green line shows the total number of patents granted.

[2] Those who subscribe to the first view often highlight 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, mismeasured in the decades that followed World War II, when many new consumer goods and technologies were introduced. Second, many implications of the growth mismeasurement thesis, such as faster productivity growth in sectors with less potential for mismeasurement, 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 fast, ICT-driven productivity growth; for example, industries with more-intensive use of ICT (outside of the ICT-producing industries themselves) have exhibited, if anything, *slower* growth of *nominal* and *real* value-added (Acemoglu et al. (2014)).

software and safety sensors, but will also require complementary improvements in energy storage, drivetrains, and tire adhesion. Consequently, when some of those innovations, say batteries, do not keep pace with the rest, we may simultaneously observe rapid technological progress in a subset of inputs and yet slow productivity growth in the aggregate. The bottleneck created by slow progress in battery technology, in this example, is *endogenous* in the sense that it is the advances of non-battery inputs that have caused batteries to become 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, current bottlenecks may offer the potential for significantly faster aggregate productivity growth: rapid progress in these technologies could enable broader gains that are held back at present.

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 power the mobile electronics and electric vehicle industries (figuratively and literally), provide a key example. Batteries were a bottleneck even prior to the 1970s, when the best available technology for rechargeable batteries (lead-acid electrochemistry) had low energy density, a slow charging rate, a short life cycle, and an unwelcome property of releasing explosive hydrogen gas while recharging. Lead-acid batteries were succeeded in the 1970s by nickel-cadmium (NiCad) 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; 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 vehicles for mass production, but also catalyzed a host of unforeseen innovations: a surge in onboard automotive processing power, enabling vehicle autonomy; battery-powered drone aircraft, now used in weather forecasting, emergency response, construction planning, filmmaking, and building inspection; and the emerging electric passenger-airplane industry. In awarding the Nobel Prize in Chemistry 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 "wireless revolution."[3]

Even more foundational to the current era is the transistor, an electronic switch that is capable of amplifying, switching, and rectifying electrical signals (Park et al., 1976). Through the 1950s, electromechanical switches and vacuum tubes were a clear bottleneck. Though used in all kinds of electronic devices, telephone lines, radios, transmitters, audio amplifiers, and early

[3]See https://www.nobelprize.org/prizes/chemistry/2019/popular-information/.

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 breaking the bottleneck that had impeded progress in technologies as disparate as computers, long distance telephones, and audio amplifiers. Due to its extraordinary switching speed, the transistor also ushered in the age of digital communications. 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 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 patenting surge in electronics and ICT would not have been feasible without this breakthrough 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 sight-lines to charted objects or a combination of optical instruments, precise clocks, and detailed tables to track progress. Traditional navigation was supplemented with radio-positioning systems in the 1970s, but these tools suffered from 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: time-keeping with atomic-level accuracy. First launched in 1978, GPS satellites now 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, it was opened to worldwide public 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. These include: precision agriculture, mining, and oil exploration; atomic-precision time information for synchronization of power transmission systems; remote surveying for geology and weather prediction; and innumerable consumer-facing services such as ride-hailing, targeted advertising, and object trackers.

We first outline a simple conceptual framework that helps formalize the ideas embodied in the examples above. In our model, technological advances (modeled as quality improvements) in a given sector depend upon simultaneous improvements in the sector's 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 that links growth in sectoral TFP to both the average TFP and the dispersion (variance) of TFP among that sector's inputs. We estimate this equation using 462 manufacturing industries between 1977 and 2007, and also for the entire US economy between 1987 and 2007 by combining our manufacturing data with 42 non-manufacturing industries.

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 doubling the variance of input-supplier TFP growth for a sector is associated with about 0.9 percentage points slower TFP growth for that 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, in an accounting sense, explain 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 1977–1987 level, then aggregate TFP growth in manufacturing would have been slightly faster (rather than considerably slower) in 1997–2007 than in either of 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 and a simultaneous increase in the TFP growth of each of the bottom 50% of industries—so as to keep average upstream TFP growth the same—would have led to 1.1 percentage points higher aggregate TFP growth in manufacturing. 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 broadly robust. They hold for the entire economy and within the manufacturing sector (where TFP is better measured), they are present in weighted and unweighted specifications, in different subperiods, with varying additional controls, and with alternative measures of productivity dispersion. We also verify that these patterns are driven by outliers and are not exclusively due to the rapid advances in computers and electronics sectors (though these sectors do play a central role in our results).

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 an external source of identification for the variance of supplier TFP growth and obtain very similar results. We also document that it is the contemporaneous dispersion of TFP among suppliers, not the future dispersion, that predicts an industry's own TFP growth.

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).[4] We show that this is unlikely to be the case. For one, we document that our results are driven by TFP, not by quantities and prices. More importantly, we document a similar relationship in patents: sectors with greater patenting variance across "idea suppliers" are less likely to patent themselves.[5] We also establish the same relationship at the firm level: firms facing greater variance of patenting activity across the patent classes that they cite are less likely to patent themselves.

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 suggestive of a potentially important linkage between (endogenous) productivity bottlenecks and productivity growth. Although further work is needed to test whether unbalanced sectoral innovation is indeed constraining aggregate productivity growth, our evidence raises the possibility of a more nuanced explanation for the productivity slowdown experienced by industrialized nations than is available in current literature—once major breakthroughs occur in sectors acting as bottlenecks, there should be an acceleration of both industry and aggregate productivity growth.

A conceptual issue raised by our paper is whether the dispersion of productivity growth across sectors is inefficiently unbalanced. High dispersion may result either from evolving technological opportunities or from inefficient allocation of research effort across industries. Our strategy is not geared toward identifying which allocation would be most efficient. Nevertheless, our evidence indicates that a more balanced trajectory of technological change would generate

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

[5]Specifically, using the Cooperative Patent Classification (CPC) scheme, we look at the mean and variance of patenting at the "upstream" patent classes. Upstream classes are constructed according to the citation network, which follows the approach in Acemoglu et al. (2016).

We do not mix the patenting and TFP analyses, both because the idea network based on citations and the input-output network are different and because 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.

substantial aggregate gains.

Our paper is related to a small but growing literature on the causes of the productivity slowdown. Alongside the views that productivity growth is high but mismeasured or, alternatively, that good ideas are becoming increasingly scarce, several other perspectives may help to explain the productivity slowdown.[6] First, and most closely related to our work, several authors have argued that productivity growth from new technologies, especially from new general purpose technologies, tends to lag the underlying breakthroughs substantially because the relevant 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. (2021) 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 emphasizes how the extent and duration of the productivity slowdown depend on the sectoral imbalance of innovation and the speed with which breakthroughs can take place in lagging sectors—rather than just slow adjustment in general purpose technologies.

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

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

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 the advancement of (and the balance across) the set of exogenously specified inputs. Our framework also relates to the motivating model in Acemoglu et al. (2016), where patenting activity in a sector depends on the number of patents in "upstream" sectors that the given sector typically cites, and the more detailed investigation of differential knowledge flows over the ideas/citation network in the recent work by Liu and Ma (2021). The key distinction between our approach and prior work is our focus on the drag that dispersion across sectors imposes 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 (e.g., Long and Plosser (1983), Acemoglu et al. (2012), and especially Acemoglu and Azar (2020)) and also from canonical quality-ladder models (e.g., Aghion and Howitt (1992), Grossman and Helpman (1991)).

Suppose 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.[7] In addition, each sector uses labor, L_{it}, and inputs that are necessary for production, X_{ijt}, which are those in the time-invariant set S_i.[8] 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 labor in production.

We model technological improvements by using a quality-ladder structure. In particular, we assume 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 of λ.

Our critical assumption is that the arrival rate of innovations depends on the distribution of input technologies that 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 innovations at time t, h and H are monotone continuous functions, and we normalize $H(0) = 0$.[9] Different choices for these functions give different relationships between the distribution of a sector's input quality and its innovation propensity. For example, we could take $h(x) = x^\rho$ and $H(x) = x^{1/\rho}$ to obtain a constant elasticity of substitution (CES) aggregator. Particularly important in this context is whether the function h in equation (2) is convex or concave. The former indicates that innovation in each sector is determined by its most advanced inputs, which means that innovations across input sectors are substitutes, implying that greater (mean-preserving) dispersion of technological know-how across inputs helps innovation. Alternatively, the concave case arises when innovations across different input sectors are complements, so that greater (mean-preserving) dispersion hinders innovation. We consider the concave case to be empirically relevant because it captures the intuitive idea, highlighted by the case studies in the Introduction, that new product and quality

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

[8] 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.

[9] We have equated the importance of an input 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 that we use in our empirical work. We also consider an alternative where the importance of an input innovation is measured by the number of citations to the innovation by patents from other industries.

improvements require simultaneous improvements in a range of inputs, and that if some of the relevant inputs fall behind, they will act as a bottleneck, slowing technological progress.[10] In both the convex and concave cases, because 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:

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

where $\bar{A}_{it} \equiv \sum_{j \in S_i} \alpha_{ij} A_{jt}$ is the (cost-share weighted) mean of the productivities of the inputs to sector i, and $\mathrm{var}(\{\alpha_{ij} A_{jt}\}_{j \in S_i})$ is the (weighted) variance of those productivities. 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:

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

where $\eta^i_{mean} \equiv H'(0)h'(\bar{A}_{it})$ represents the effect of the mean productivity 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. 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 the imbalance of innovations across key input sectors in the economy may hold down aggregate productivity growth.[11]

To illustrate this point succinctly, suppose that $S_i = S$ for all i and some $S \subset \{1, \ldots, N\}$ and that $\alpha_{ij} = \alpha_j$ 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 the weighted variance, $\mathrm{var}(\{\alpha_{ij} A_{jt}\}_{j \in S_i})$, is given by σ^2. Equation (3) implies 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.[12]

[10] The inputs that need to make technological advances before sector i can successfully innovate may be a subset of the inputs in S_i. Because we do not have a way to empirically determine which subset of inputs is important for innovation, we assume that all inputs in S_i are relevant, then verify robustness using other measures of industry linkages.

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

[12] As we discuss 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 every sector is the same.

2.2 Endogenous Innovation Effort

It is straightforward to endogenize innovation and characterize the general equilibrium.[13] While endogenous 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 efforts. 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 simultaneously pursue similar ideas. We include $1/\gamma$ as a constant in front of the H function, for simplicity. Note also that the H function here represents a pure knowledge externality, and thus the fact that sector i builds on the industries in the set S_i does not generate additional profits for these industries.

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 and may also include sector-specific distortions, as well as misperceptions or fads among researchers (i.e., motivations of researchers 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 here also, fads and misperceptions may affect rewards as well).

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

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

and thus

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

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

[13] Finding the general equilibrium will also require asked to solve for the wage rate (and the allocation of labor across sectors) and the interest rate (as a function of the aggregate growth rate of the economy). We do not derive these (standard) aspects of the general equilibrium.

Equation (5) emphasizes that, if the cost of research, κ_i, varies 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 technological progress rates across sectors would 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 the mean productivity of new innovations, thereby improving aggregate productivity (and welfare). Conversely, if differences in κ_i 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 Industry Database.[14] 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 prices. We include 462 manufacturing industries, corresponding to six-digit NAICS codes. In accordance with the literature, TFP is defined as the residual 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 42 non-manufacturing industries, corresponding to three-digit NAICS codes, from the BLS Major Sector and Major Industry Total Factor Productivity database from 1987–2011. As with the manufacturing data, TFP outside of 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.[15]

We construct input-output tables using the detailed Make and Use tables provided by the US Bureau of Economic Analysis from 1977–2007, which are available every five years, corresponding to the years of the Economic Census. These tables provide information on the amount

[14]https://www.nber.org/research/data/nber-ces-manufacturing-industry-database.

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

that each industry produces of various commodities and the amount that they spend on each commodity, respectively. From these two tables, we construct our basic input-output network $\{\alpha_{ijt}\}$ whose entries are the dollar value of inputs that industry i uses from industry j at time t relative to the dollar value of its total intermediate costs. Since each year's release of these tables uses industry coding 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 in the Appendix. Appendix Table A1 presents summary statistics across upstream (supplying) and downstream (customer) industries. Panel A shows results for manufacturing industries only, during 1977–2007, while Panel B depicts averages for all industries, over 1987–2007. In the former, we see that the average five-year TFP growth across manufacturing sectors was 1.8 percentage points. The average TFP growth of upstream manufacturing industries is substantially higher, at 3.3 percentage points, reflecting the fact that more-productive industries are used more intensely as intermediate inputs.

To explore innovation outcomes directly, we look at patent data, starting from the UCB Fung Institute Patent Data Project, which spans the years 1976–2016. These data include every patent application and patent granted by the USPTO during this time period. While the data do not include patents granted outside the United States, they contain patents filed at the USPTO by non-US firms. The data include patent classification codes, application dates, and (importantly) cross-citations to other patents. Firm names and locations are cleaned using machine learning and natural language processing (see Balsmeier et al. (2018) for additional details on the disambiguation algorithm). The patent classification codes refer to 633 unique Cooperative Patent Classification (CPC) classes. We construct a time series that tracks the total number of patents in each CPC class by application date, as well as a similar time series for the patenting activity of each firm.[16]

We also use these data to construct the CPC-level citation network (what Acemoglu et al. (2016) refers to as the "innovation network"), which represents the knowledge flows between CPC classes. Specifically, following Acemoglu et al. (2016), we calculate a citation network, γ_{cj}, whose entries are the fraction of citations to patents in CPC class j among total citations of patents in CPC class c. To achieve greater precision and remove the time-dependent measurement error problem introduced by the increasing number of patents over time, we use the average number of citations for each class over the entire sample. We exclude all within-CPC citations, meaning citations by patents in CPC c to other patents within CPC c.

Likewise, we construct firm-level citation networks. In this case, we calculate a citation

[16]If there is more than one CPC code provided for a patent, we use the first-reported (i.e., primary) CPC code.

network, ω_{kc}, whose entries are the share of citations by firm-k patents (i.e., patents that belong to firm k) to the patents of other firms within the CPC class, c. We exclude all within-firm citations (i.e., citations by firm-k patents to other firm-k patents).

Lastly, we supplement the domestic, US data with data for select European countries. We use data on value-added and TFP from the 2012 EU KLEMS Growth and Productivity Accounts. In this exercise, we use data from 1987–2007 for 30 industries in Austria, Finland, France, Germany, Italy, Netherlands, Spain, and the United Kingdom. We combine these data with country-specific input-output tables from the GGDC World Input-Output Database for 2000. The relevant entry in the world IO table, $\alpha_{ik,jl}$, is the share of inputs for industry i in country k that came from industry j in country l.[17] Panel C of Appendix Table A1 presents the TFP growth for this sample. Overall, the average five-year TFP growth within this sample was 4.6 percentage points. For patenting outside of the United States, we use Google Patents global patent data from Liu and Ma (2021), which contains patents from over 40 major patent authorities around the world. Each patent is assigned to a geographical unit using the country of residence of the inventor, the country of residence of the assignee(s), and country of the patent authority, in that order. We construct the number of patents in each country in each CPC code in each year, using the date of application for each patent. We further restrict our attention to the 20 countries with the most patenting over the sample period.

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—with this dispersion representing an imbalance under our hypothesis. Concretely, we estimate a version of equation (3), derived above, using data on 462 six-digit, NAICS-based manufacturing industries between 1977 and 2007, and 42 three-digit non-manufacturing industries in 1987–2007. We also report the quantitative implications of these estimates and document their robustness to additional controls, different sample periods, and sources of variation in productivity growth.

[17]In our baseline specification, this share includes inputs from all other countries. We explore alternative definitions and, in Appendix Table A6, show that our results are mostly driven by TFP growth patterns among a country's domestically sourced inputs, with a more limited role for imported intermediates.

4.1 Main Results

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

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

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,

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

and $\sum_j \alpha_{ijt-1} \Delta TFP_{jt}$ is the average TFP growth among the suppliers of industry i during the five-year time period, calculated using the α_{ijt-1}'s as weights. Recall that α_{ijt} represents the ratio of industry i's spending on inputs from industry j relative to its total intermediate spending time t. The variance of TFP growth among the suppliers of industry i is also computed using these cost shares as weights. In addition, $\mathbf{X}'_{it-1}$ denotes a vector of other (predetermined) covariates, which in some specifications includes sector fixed effects, introducing sector-specific linear trends; 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 technology 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), η^i_{mean} and $\eta^i_{variance}$, we have imposed constant coefficients, which should be interpreted as local average treatment effects.

Throughout, we always control for the mean effect of supplier TFP growth, but the main coefficient of interest for our study is $\beta_{variance}$, which captures the effect of supplier TFP growth dispersion (or innovation dispersion, in the case of our patent analyses) on a sector's productivity (innovation), holding constant the mean of supplier TFP growth (innovation). We expect this coefficient to be significantly negative if, as we hypothesize, imbalances in the rates of technological progress across an industry's suppliers impose a productivity penalty on the industry.

Table 1 reports estimates of equation (6) for 462 six-digit manufacturing and 42 three-digit non-manufacturing industries. Panel A is for manufacturing industries, where 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. All models include time fixed effects, while each specification includes an alternative with industry fixed effects, which allow each industry to have its own linear time trend in TFP. Odd-numbered columns include no covariates other than time dummies, while even-numbered columns 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 1–5, are unweighted. We weight industries by their share of 1987 real value-added in columns 6–8.

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 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 our baseline specification, column 2 of Panel A, the coefficient estimate of the variance term is -0.744 (standard error = 0.121). Adding linear industry trends in column 3 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 and only slightly larger. Figure 3 depicts the industry-level variation that produces 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 downstream industry TFP growth, and the right panel showcases the strong negative relationship between the variance of supplier TFP growth and downstream industry TFP growth.

The specification in (6) is a natural one, using the variance term to capture the effects from supplier TFP growth dispersion, as in our second-order approximation above. Nevertheless, it is useful to see whether well-performing and poorly performing supplier sectors both impact TFP growth. To investigate this question, columns 4 and 5 replace the variance term with TFP growth in the 10^{th} and 90^{th} percentiles of the (weighted) TFP distribution of suppliers (while we continue to control for mean supplier TFP growth). Consistent with our hypothesis, holding mean supplier TFP fixed, higher 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 these relationships typically exhibiting statistical significance at the 5% level or below). Lastly, columns 6 through 8 replicate the main specifications, but with shares of real value-added in 1987 used as industry weights. These weighted estimates are very similar to the unweighted specifications.

Overall, the estimates in Table 1 uniformly show a negative estimated impact of TFP growth dispersion across a sector's suppliers on own-industry TFP growth. In terms of our motivating conceptual framework, this suggests that productivity growth in a sector is held back when advances among its suppliers are unbalanced. In the rest of the paper, we demonstrate the robustness of these results and document a similar relationship in innovation activity. Before moving on to innovation, we draw out the quantitative implications of the productivity growth estimates in the next subsection.

4.2 Quantitative Implications

The results in Table 1 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 United States. For this to be the case, two conditions must be satisfied. First, the coefficient estimates in Table 1 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. Panel A of Figure 4 depicts the simple variance of TFP growth across all manufacturing industries, while Panels B and C show the average variance of industry supplier TFP growth: for manufacturing only and for the economy overall, respectively. Both within upstream manufacturing and across all manufacturing industries, 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 mid-1970s and now is three times as large, around 0.006. As suggested by the patenting time-series in Figure 1, the electronics and computer sector accounts for a large portion of the increase in TFP variance through the 1990s. The right-hand side plots of Figure 4 document that when this sector is taken out, the rise in the variance of TFP growth is noticeably smaller—though still present—in recent decades. When we zoom out to include non-manufacturing supplier industries, there is a similarly large increase in the variance of TFP growth from the 1980s 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 be explained by the rising variance of TFP growth? Figure 5 addresses this question by applying the value-added weighted estimates reported in column 7 of Panels A and B in Table 1. We find a sizable productivity penalty from TFP growth dispersion. The estimates imply that TFP dispersion reduced manufacturing TFP growth by over 0.5 percentage points in both of the recent 10-year periods (as shown by the green bars in the left panel) and TFP growth in the entire economy by about 0.2–0.3 percentage points during each five-year period over 1992–2007 (as shown by the green bars in the right panel). 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 between 1997–2007 had remained at its average level from 1977–1987, our estimates imply that aggregate TFP growth in manufacturing between 1997–2007 would have been 3.1% (as shown by the orange bars in the left panel) instead of the observed TFP growth of -1.6% (as shown by the blue bars). This counterfactual growth rate is slightly more rapid than the 3% actual TFP growth in manufacturing during 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 1 can potentially account for the bulk of the US productivity slowdown in recent decades. We stress that these magnitudes are suggestive but far from definitive, given the limitations of our measurement and identification.

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 2 lists illustrative examples of the fastest-growing industries, which are defined as those that have had the largest impact on supplier TFP variance between 1997 and 2007. These industries 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 the 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 bottom 50% of industries just enough to keep the average TFP growth constant.[18] In this scenario, the variance of supplier TFP growth between 1977 and 2007 would have been 40% lower and aggregate TFP

[18]Like the fastest-growing industries, the bottom 50% industries are defined in terms of their contribution to supplier TFP variance between 1997 and 2007. Appendix Table A2 reports the full set of industries corresponding to each panel of Table 2.

growth in manufacturing would have been 1.1 percentage points higher.

The remaining panels of Table 2 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 period. These include pharmaceutical preparation, basic organic chemicals, printed circuit assembly, and turbine generators. Panel C reports example industries that are most *bottlenecked*—that is, held back by the uneven innovation across their suppliers. These include surgical and medical instruments, gas engines, and industrial valves.

4.3 Endogeneity Concerns

Since the estimates in Table 1 are obtained from regressions of an industry's 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. In this subsection, we explore two strategies that, in net, lend support to the case that these results are informative about the effects of productivity bottlenecks.

Our first strategy is to isolate industry productivity changes that emanate from common technological developments across several advanced economies. We do this in Table 3 by exploiting changes in industry TFP in major OECD countries, as reported by the 2012 EU KLEMS Growth and Productivity Accounts. For this exercise, we focus on all 504 industries (both manufacturing and non-manufacturing), mapped to 29 EU KLEMS industries.[19] In Panel A, we use the mean and variance of supplier TFP in France, Germany, and United Kingdom as instruments for the corresponding variables in the United States. To purge measurement error in these instruments, Panel B uses the rank of TFP growth by industry within country. In both panels, columns 1 and 2 present the baseline OLS results, while columns 3 and 4 depict 2SLS estimates.

The first-stage F-statistics are given at the bottom of Panels A and B in Table 3; these are somewhat low in Panel A, but higher in Panel B, where we use TFP ranks as the instrument (the full first stages are reported in Appendix Table A5). This motivates the Limited Information Maximum Likelihood (LIML) estimates presented in columns 5 and 6, which are consistent even in the presence of weak instruments. These estimates confirm that our findings are not driven by weak instruments.

[19]More specifically, we calculate these instruments using the US-based IO table, but taking the TFP growth across industries from each of three European countries (France, Germany, and UK). Because the international industry data are more-aggregated than our underlying NAICS data, six-digit NAICS codes are mapped to the most-similar international industry code available, the TFP growth value observed in the European instrument panel is assigned to US industries based on these mappings. To reflect this, we cluster the standard errors at the level of the 29 KLEMS industries in Table 3.

The IV estimates of the relationship between industry TFP growth and supplier TFP mean and variance correspond closely to our earlier OLS estimates. In both panels, the OLS and IV estimates are very similar across columns 1–2 and 3–4. For example, in columns 1 and 3 of Panel A, which do not include industry fixed effects, the OLS coefficient on the variance term is -0.876 (standard error = 0.155) and the IV estimate for the same variable in the same specification is -0.902 (standard error = 0.385). The variance term estimates are also quite close in columns 2 and 4, where we add industry fixed effects. We see a similar pattern in Panel B when we exploit the variation in the rank of TFP growth: -0.667 (standard error = 0.445) for the IV estimate without industry fixed effects in column 3 and -1.480 (standard error = 0.661) for the IV estimate with industry fixed effects in column 4. The LIML estimates in columns 5 and 6 are also comparable. For example, in Panel B, the variance term's coefficient estimate is -1.482 (standard error = 0.665) with industry fixed effects—similar to the 2SLS estimates in column 4 and the OLS estimates in column 2.

The congruence between the baseline OLS estimates and the IV estimates that exploit TFP changes in other leading economies bolsters our confidence that these results are not driven by shocks that are common across US industries and their suppliers. We also note that since the IV coefficient estimates are similar to the OLS estimates from Table 1, the implied quantitative magnitudes are comparable as well.

Panel A of Table 4 explores whether there is a correlation between *future* average and variance of TFP across suppliers and an industry's *current* TFP growth. Such a relationship would be concerning for the interpretation of productivity bottlenecks as a constraint on TFP growth. Across the eight columns (the first four for manufacturing industries and the last four for all industries), we do not see any evidence that future variance of TFP growth of suppliers has a negative relationship with current TFP growth of an industry. When we focus on all industries and look at the relationship between future variance and current TFP growth, the correlation is positive, but it disappears when we include our main regressors, contemporaneous average and variance of TFP across suppliers. It is also zero across all specifications for manufacturing industries. This pattern is reassuring for our overall interpretation.

Panel B of the same table explores whether a similar relationship exists between the variance of TFP growth among an industry's customers and its own TFP growth. Such a correlation is a distinct possibility, since many industries have customers and suppliers that are overlapping. The general pattern is that customer variance is also negatively correlated with an industry's TFP growth, but typically only when it is entered by itself (without the variance of TFP among input suppliers). When both sets of variables are included, the coefficient on customer variance becomes less significant and smaller, while the variance of TFP across input suppliers remains negative and significant. This pattern is broadly supportive of our overall interpretation, even if it raises the possibility that in some specifications there is a high enough correlation between downstream and upstream variances that we cannot rule out additional effects working from customers' TFP growth.

Panel C of Table 4 investigates whether mean-reversion dynamics may be confounding our estimates. In particular, if TFP growth is serially correlated, then failing to account for this could lead to a spurious relationship between supplier TFP variance and own TFP growth. Panel C addresses this issue by including the lag of own TFP growth, as well as lagged input average and variance terms (in some specifications). Overall the results prove quite robust, and the statistical significance and quantitative impact of the input variance term are hardly affected. For example, in specifications that include only the lagged dependent variable, the input variance has a coefficient of -0.747 (standard error = 0.115) for manufacturing industries and -0.923 (standard error = 0.152) for all industries, similar to our baseline findings in Table 1 in both cases. The estimates are again similar when we include lagged input average and variance terms.

4.4 Robustness

Table 5 further investigates the robustness of our results to a battery of controls and specifications. For brevity, we focus on the manufacturing sample and report analogous results for all industries in the Appendix (see Appendix Table A4). Panel A documents robustness for the specification without industry fixed effects, while Panel B includes industry fixed effects that allow for linear trends in industry TFP.

For ease of reference, column 1 reports our baseline specification, which includes only the mean and the variance of TFP as well as time dummies. In column 2, we show analogous results using a weighted specification (weighting industries by real value-added). This specification is particularly useful, since our quantitative explorations aggregate across industries using real value-added weights. The estimates are very similar to those in column 1.

Column 3 estimates the same models but now using 10-year periods rather than the five-year intervals in Table 1. This specification purges higher-frequency variation in TFP and focuses on longer-term variation. 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 following rapid growth in the recent past, this would not correspond to an imbalance but rather to a potential rebalancing. Column 4 checks this possibility by adding the covariance between the supplier TFP growth in the current and the prior periods to our specification.[20] Intuitively, this covariance term accounts for potential persistence and reversal patterns in industry-level TFP changes. We find that the covariance of TFP across periods does not meaningfully affect the relationship of primary interest. The coefficient on the variance term in Panel A is only slightly larger, -0.640 (standard error = 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 the coefficient on the variance of upstream TFP growth remains unaffected by the inclusion of this 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 5 provide additional robustness checks. Another factor that could affect measured industry TFP is changing import penetration. Column 5 controls for average imports from China by other countries in the industry and in the input-weighted average of the supplying industries (from Autor et al. (2013)), addressing the concern that Chinese import penetration may itself impact productivity growth (e.g., Autor et al. (2020)). Accounting for imports does not appreciably change the coefficient on supplier TFP variance.[21]

We noted the importance of the electronics and computer sectors above. Column 6 confirms that the negative relationship between industry TFP growth and supplier TFP dispersion holds even when computers and electronics manufacturing (NAICS 334) is excluded from the estimation sample as well as from the calculation of upstream metrics. With these key sectors excluded, the variance term is less precisely estimated, as expected. Nevertheless, it remains statistically significant at the 5% level or below in all of our specifications: -1.231 (standard error = 0.587) in Panel A and -1.197 (standard error = 0.624) in Panel B. These estimates reveal that that our hypothesized mechanism is present even when the ICT and electronics sectors are excluded, but also that the ICT and electronics sectors showcase our mechanism and contribute substantially to its identification and quantitative implications (as corroborated by the examples in Table 2).

Column 7 shows a similar relationship to our baseline results when we estimate a robust regression that downweights outliers that have a major effect on the slope of the relationship between upstream TFP variance and downstream TFP growth. Notice that in this case, the standard error of the input variance term is much smaller, highlighting that outliers were,

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

[21]An additional concern is whether our variance term is misspecified, because it does not account for the productivity of offshored and imported inputs. Here we note that this concern would create attenuation towards zero, and hence is unlikely to account for our findings. It also does not apply when we turn to patenting, since our analysis there will include foreign patents as well. We further discuss this issue in Section 6.

indeed, reducing the precision of our estimates, though not affecting their magnitudes much.

Column 8 confirms that the results are again similar when we use a fixed input-output matrix, rather than the time-varying input-output matrix from our baseline specification. Column 9 probes robustness to our definition of input shares. Here, we define upstream shares, α_{ijt-1}, as total-cost shares rather than as intermediate-cost shares (as in our baseline specification). These two share measures will differ to the extent that the intermediate share of total costs varies across sectors. The results are once again very similar. Finally, column 10 excludes own three-digit industry when constructing the input-output network. This does change the magnitude of the coefficient estimates but not their signs or statistical significance. Appendix Table A4 shows analogous estimates for the entire economy, rather than just the manufacturing sector. These results are again similar to our baseline estimates.

Since our empirical analysis is confined to the 462 manufacturing (or the 504 total) industries, our estimates will not capture any imbalances in innovation or productivity growth that happens at more-disaggregated levels. To explore whether these more-micro imbalances may also matter, and to further probe the robustness of our results, in the Appendix we use estimates of within-industry, across-establishment TFP growth from the US Census Bureau's Dispersion Statistics on Productivity (DiSP). These measures of dispersion have also increased during our sample period, but Appendix Table A3 documents that the average upstream TFP growth dispersion among input suppliers, when added to our regression, is not statistically significant and does not change the relationship between our measure of supplier TFP growth dispersion and own TFP growth.

In summary, these results 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.

4.5 Prices, Quantities and Productivity

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

We investigate the role of these neoclassical channels in Table 6, where we add the mean and variance of supplier prices and employment levels to our baseline regressions. For comparison, columns 1 and 2 restate our baseline estimates. The alternative specifications in columns 3–8 indicate that controlling for these neoclassical channels does not qualitatively change the relationship between supplier TFP variance and industry TFP growth (and that the coefficients on these channel variables are typically insignificant). For example, when we include the mean and variance of supplier *prices* in column 3 (without industry fixed effects), the TFP variance term has a coefficient of -0.686 (standard error = 0.232), which is 90% of the baseline estimate, though less-precisely estimated. When we include the mean and variance of supplier *employment* levels in column 5 (also without industry fixed effects), the coefficient estimate on TFP variance is -0.703 (standard error = 0.118), which is nearly identical to the baseline estimates in column 1. The results remain similar when we include *both* sets of variables (prices and employment) together. When we include industry fixed effects, the estimates are once again similar to our baseline results.

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

5 Innovation

This section investigates whether innovation, as encoded in patents, is one of the underlying mechanisms that could explain our results. For this exercise, we replace the input-output network (comprised of α_{ijt} entries) with the patent citation network (corresponding to the γ_{cj}'s capturing citation patterns across CPCs). Our sectoral analysis starts at the CPC level, but we also consider firm-level results later in this section. The main question explored in this section is whether a greater imbalance of innovation across upstream sectors or firms ("idea suppliers") reduces the innovation of a downstream sector or firm. We will see that the answer to this question is a strong *yes*.[22]

[22]Because the mappings between CPCs and SIC/NAICS classifications are imperfect, we do not explore the relationship between upstream patenting and downstream productivity growth.

5.1 CPC-Level Results

We begin the analysis at the patent-class level and estimate the following variant of equation (6):

$$\Delta \text{Patent}_{ct} = \beta_{mean} \sum_{j} \gamma_{cj} \Delta \text{Patent}_{jt} + \beta_{variance} \text{VAR} \left(\Delta \text{Patent}_{jt} \right) + \delta_t + \varepsilon_{ct}, \tag{7}$$

where t refers to five-year time periods, ΔPatent_{ct} is a measure of patenting growth within CPC c during the five-year time interval denoted by t,

$$\text{VAR} \left(\Delta \text{Patent}_{jt} \right) \equiv \sum_{j} \gamma_{cj} \left(\Delta \text{Patent}_{jt} - \sum_{j} \gamma_{cj} \Delta \text{Patent}_{jt} \right)^{2},$$

and $\sum_j \gamma_{cj} \Delta\text{Patent}_{jt}$ is the average patent growth during the five-year time period among the CPCs that are upstream to c (that is, among the CPC codes that c, the focal CPC, cites). As indicated above, its entries, the γ_{cj}'s, are the share of total citations over the entire sample period from patents in CPC c that go to patents in CPC j. The upstream variance of patenting growth is also computed analogously to the upstream variance of supplier TFP, though now using the γ_{cj}'s as weights.

Table 7 presents our main estimates of the patent-based version of equation 6. The first three columns measure innovation activity by log patents, which implies that sectors with zero patenting activity are dropped (this produces a sample of around 4,326 observations for our main specifications). Columns 4–6 instead focus on the Davis-Haltiwanger-Schuh (DHS) transformation (Davis et al., 1998), which allows us to define growth rates when there are zero patents in a CPC in either the beginning or the end period.[23] This expands our sample slightly to 4379 CPC observations.

Columns 1 and 4 show a strong positive association between average patenting activity in a sector's upstream CPCs and the sector's own patenting activity. Columns 2–3 and 5–6 add the variance of patenting activity in the upstream CPCs to proxy for the imbalance of innovation activity across sectors. The latter two columns also include CPC fixed effects, which allow for linear trends at the CPC level. The estimates show a powerful negative effect of upstream variance. In column 2, for example, the coefficient estimate for the variance is -0.959 (standard error = 0.266). The variance estimate remains essentially unchanged in column 3 when CPC fixed effects are included. The coefficient estimates are very similar in columns 5 and 6 with the DHS transformation, though standard errors are somewhat larger.

Table 8 is the patenting analogue of Table 3 from our IV analysis for TFP, but now focusing

[23] The DHS transformation for a variable X is $\frac{X_t - X_{t-1}}{\frac{1}{2}(X_t + X_{t-1})}$.

on patents and exploiting variation in upstream patenting among foreign patents contained within the Google Patents global database. The estimates in Table 8 are broadly supportive of the negative relationship between upstream variance and a sector's own patenting. Panel A depicts specifications using the change in log patenting (as in columns 1–3 of Table 7), while Panel B shows results with the DHS transformation (as in columns 4–6 of Table 7). In each panel, columns 1 and 2 show the OLS relationship in this sample, columns 3 and 4 report 2SLS estimates, and columns 5 and 6 present the LIML estimates (which are again motivated by the weak first stages in columns 3 and 4). Across essentially all columns, we see negative and statistically significant estimates of the impact of upstream variance.

Following the design of Table 4 above, Table 9 explores whether *downstream* patenting variance also matters; the possible relationship between future upstream variance and current patenting; and whether mean-reversion dynamics may be confounding our results. Reassuringly, the results in this table confirm the robustness of the estimates in Table 7 to these checks. In particular, in Panel A future citations have a smaller and often insignificant coefficient when entered at the same time as our main citation variables, while the coefficient on our citation variance measure remains similar to the baseline estimates.

In Panel B, downstream variance—that is, variance among *citing*, rather than *cited*, patents—is negative and significant when entered by itself, which reflects the fact that, just as in the input-output network, upstream and downstream measures are correlated. Nevertheless, when we also include our upstream citation variables, the downstream variance is no longer statistically significant and is in fact positive in most specifications, while our upstream citation variance has a similar coefficient to our baseline estimates and is statistically significant with log patents in columns 2 and 4, though it becomes less precise with DHS in columns 6 and 8.

Finally, in Panel C we find that the inclusion of lagged patenting (the dependent variable) and the lagged citation average and variance terms has very little effect on our results when focusing on the log patents measure. This is also the parent we find with the DHS transformation without fixed effects (column 6), though with DHS and CPC fixed effects, the coefficient on the citation variance falls and becomes statistically insignificant (column 8).

Table 10 confirms the robustness of our main CPC-level estimates to the same battery of tests we conducted in Table 7 for TFP. (For brevity, we focus on log patents as the dependent variable.) We see broadly similar patterns across specifications that: include or exclude CPC trends; are weighted by their share of total patenting; are at the 10-year frequency; include the covariance term; exclude the ICT and electronics sectors; leave out all citations to patents in the focal sector's three-digit CPC; limit the sample to years before 2005 in order to make the sample more similar to the data used for the TFP growth analyses; or focus on patents filed by US residents. The only two specifications in which the variance term is significantly weakened are columns 5 and 8. The former of these excludes the ICT and electronics sectors, and the weaker results likely reflect the factors discussed for the analogous specification in Table 5—computers and electronics are emblematic of the imbalances that are our focus, so excluding these industries weakens the relevant economic forces and the precision of the estimates. The latter, column 8, excludes the 54% of total patents filed non-US residents at the USPTO, which likely accounts for the reduced precision of these estimates.

Quantitatively, these estimates suggest that upstream innovation imbalances have a major impact on overall innovation. For example, the weighted coefficient estimate in column 2 of Panel A in Table 10 suggests that a one-standard deviation higher upstream variance (which is 0.03) is associated with a decline in the growth rate of patenting in a CPC code of 0.042 log points. This is a 47% reduction relative to the weighted mean of patenting across sectors, which is equal to 0.09 (weighting by the total number of patents in the CPC code in the initial five-year period). These numbers are in the same ballpark as those implied by our TFP models.[24]

In sum, although the results in this subsection show a few specifications where the estimates are less stable than our main results reported in the prior section, they are overall supportive of a robust negative association between upstream variance of innovation activity and downstream patenting at the CPC level.

5.2 Firm-Level Evidence

We next turn to the firm-level relationship between upstream imbalances and patenting. For this exercise, we disaggregate the patents data to the firm level and allow for variation across firms: specifically, the extent to which they rely on different CPCs for their patenting. This produces our firm-level citation network, summarized by $\{\omega_{kc}\}$, representing citations by firm k to CPC class c. We estimate the following equation:

$$\Delta\text{Patent}_{kt} = \beta_{mean}\sum_{c}\omega_{kc}\Delta\text{Patent}_{ct} + \beta_{variance}\text{VAR}\left(\Delta\text{Patent}_{ct}\right) + \delta_t + \varepsilon_{kt}, \tag{8}$$

where t refers to five-year time periods, ΔPatent_{kt} is a measure of patenting growth of firm k during the five-year time interval denoted by t,

$$\text{VAR}\left(\Delta\text{Patent}_{ct}\right) \equiv \sum_{c}\omega_{kc}\left(\Delta\text{Patent}_{ct} - \sum_{c}\omega_{kc}\Delta\text{Patent}_{ct}\right)^2,$$

[24] In particular, our main TFP estimates from column 2 of Table 1 suggest that a one-standard deviation increase in upstream TFP growth variance is associated with TFP growth that is 0.035 percentage points lower.

and $\sum_c \omega_{kc} \Delta \text{Patent}_{ct}$ is the average patent growth in the five-year time period among the CPCs upstream to firm k (meaning those cited-to by firm k). As indicated above, these are calculated using the share of total citations over the entire sample period by firm k's patents to patents in CPC c. The variance of patent growth among the cited CPCs is computed using the ω_{kc}'s as weights.

This disaggregation produces a much larger sample, consisting of almost 2 million observations at the firm level. For many firm-period combinations, however, there are no patents. Thus, in this table, we use the DHS transformation. In particular, in columns 1-3, we use the standard DHS transformation, where observations are dropped when there are two consecutive zeros. In columns 4-6, we use a modified DHS transformation, where in such cases, the transformation imputes a value of zero.[25]

Table 11 presents the main results from this exercise. The firm-level structure of the data in this table enables us to control for firm fixed effects or for CPC-times-year fixed effects, thus purging a large fraction of the variation in patenting between firms. The general pattern is a negative relationship between upstream variance at the CPC level and a firm's own propensity to patent. For example, in column 1, the coefficient estimate of the citation variance is -0.264 (standard error = 0.042). In column 3, when we include CPC-times-year fixed effects, the coefficient increases slightly, to -0.292 (standard error = 0.045). The exception to this pattern is in column 2, where we see a positive and significant coefficient when firm fixed effects are included with the standard DHS transformation. We suspect that this is driven by firms that have many zeros and thus many missing observations. Indeed, in column 4-6, when we use the modified DHS so that all zeros are kept, the coefficient on the variances are more stable and always negative (and strongly statistically significant except in column 5).

Table 12 provides a number of robustness checks for these firm-level results, considering analogous specifications to those we presented for the TFP and CPC-level patenting models and focusing on the standard DHS measure. Panel A of this table corresponds to column 1 of Table 11, while Panel B adds CPC-times-year fixed effects, as in column 3 of that table. The results are robust across specifications that are weighted by the firm's share of total patenting; control for the lagged dependent variable; change the sample period; or focus only on domestic patents. As a further robustness test, column 6 adds the mean and variance of *future* citations in a firm's patenting network. Future citations as well as our main measures are now statistically significant. Given the high degree of serial correlation in patenting within classes, these patterns are not surprising. They highlight however that future tests of our proposed mechanism should

[25]Recall from footnote 23 that the DHS transformation is $\frac{X_t - X_{t-1}}{\frac{1}{2}(X_t + X_{t-1})}$. This is undefined when both X_t and X_{t-1} are equal to zero. In the modified DHS, rather than dropping such observations, we set $\frac{0}{0} = 0$.

attempt to exploit shocks that impact patenting during certain discrete periods.

Quantitatively, these estimates imply that upstream firm-level imbalances have similarly sized innovation effects as we measure at the CPC level. For example, the coefficient estimate in column 2 of Panel A of Table 12, which shows the weighted regression specification, suggests that a one-standard deviation higher upstream variance (which is again 0.03, as in the CPC case in the previous section) is associated with a decline in firm-level patenting of 0.13 log points. This is a sizable, 73%, decline relative to a baseline of 0.18. These numbers are similar when we include CPC by year fixed effects in the same column of Panel B. Overall, these numbers are broadly comparable to those from the CPC-level analysis in Section 5.1.

6 International Evidence

Our primary analysis focuses on TFP growth and innovation in the United States (except when instrumenting domestic TFP growth and innovation with contemporaneous, foreign development in Tables 3 and 8). We supplement this evidence here by estimating a variant of equation (6) for TFP growth across European countries. As outlined in Section 3, we use the GGDC World Input-Output Database (WIOD) to construct consistent input-output linkages for 30 industries in Austria, Finland, France, Germany, Italy, Netherlands, Spain, the United Kingdom, and the United States. We fix the global input-output table at the year 2000 and focus on industry TFP growth in this cross-country sample between 1987 to 2007. These data enable us to include international input-output linkages, which we exploit in our calculations of the mean and variance of supplier TFP growth.[26]

We report these cross-country estimates in Table 13. We report the baseline specifications in the first four columns. These specifications are unweighted and include combinations of country effects, year effects, year-by-country effects, and year-by-industry effects, as noted at the bottom of each column. In column 1, we focus on a specification containing country and year effects. This estimate verifies 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 estimate, at -0.824 (standard error = 0.212).

Subsequent columns probe the robustness of this finding. Column 3 includes industry-by-year effects, so that the identifying variation is within-industry rather than cross-industry, as in the main specifications of the paper. The relationship is similar to column 2, although

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

somewhat smaller. In particular, the coefficient on the variance term is -0.535 (standard error = 0.143). Column 4 includes both industry-by-year and country-by-year interactions, restricting to variation within-industry and within-country. 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 the 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.[27] In column 9, we show that the estimates are similar when we only use each country's domestic input-output network, rather than the full international input-output table (which incorporates inputs from each country-industry pair).[28] Finally, in Appendix Table A6, we report cross-country regressions that include both own-country (domestic) values and foreign-country average values of the mean and variance of upstream (supplier) TFP growth as explanatory variables for sectoral productivity growth. These models show that own-country supplier TFP values are a far more robust predictor of sectoral productivity growth than the corresponding other-country values. This is especially the case for the variance term, where the own-country coefficient is negative and significant in all columns whereas the other-country measure is neither significant nor consistently signed. This pattern is reassuring against the concern that our upstream TFP variance terms may be misspecified because they do not include productivity growth among important intermediates (see footnote 21).

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 United States, 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 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 during 1992–2007 by 30% and 60%, respectively.

We also implement a similar specification where the outcome is patenting among firms located in different countries. We use the same CPC-citation linkages for each country (calcu-

[27] See Appendix Table A7 for robustness of our US results, aggregated to the 30 industries used in Table 13.

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

lated using all patents within the USPTO database), but we apply a variant of equation (6) for patenting growth of firms located in the 20 largest countries within Google global patent data. Table 14 reports these cross-country estimates. Columns 1 through 4 use log patents while columns 5 through 8 use the DHS transformation. Exploiting the cross-country variation, we see in columns 2 and 6 that, with either specification, there is a negative effect of upstream imbalances on patenting. This remains true in columns 3 and 7, where we include year-by-CPC fixed effects, thus identifying the relationship exclusively from cross-country variation in the upstream variance. The relationship is also broadly robust, though a little less precisely estimated, in columns 4 and 8, where we further include year-by-country fixed effects, thus focusing entirely on within-country variation.

7 Conclusion

Despite the exponential pace of innovation in the information and communications technology (ICT) and electronics sectors, aggregate productivity growth in the United States and many other industrialized nations has been disappointing since the 1970s—and only more-so since the early 2000s. Some have interpreted this pattern, variously, as reflecting a severe underestimation of quality and actual productivity growth; a temporary lull that proceeds a major surge in productivity; or an exhaustion of the potential supply of truly transformative innovations—leading to a long-term deceleration of productivity growth.

We proposed an alternative hypothesis that implies neither a permanent slowdown in productivity growth nor an incipient surge. We then investigated this new hypothesis empirically. The foundational idea of our approach is that innovation in any one industry relies on complementary innovations in—and subsequent productivity gains from—its input and idea suppliers. When innovation is unbalanced across industries, this holds back aggregate productivity growth by creating innovation "bottlenecks" along the input-output or patent citation (idea) networks.

After presenting a simple version of this productivity bottleneck hypothesis, we explored 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: an industry's productivity growth is augmented by the mean productivity growth of its suppliers (measured by TFP or innovation) and, crucially, it is hampered by the variance of their productivity growth.

Our primary evidence exploits input-output linkages and TFP growth to document the sensitivity of industry productivity growth to the mean and variance of supplier productivity growth. We supplemented this evidence by looking at patenting as a direct measure of innovation. This analysis suggests that there is a similarly powerful linkage between the innovativeness of a sector or firm and the imbalances it faces across its upstream (idea-supplier) sectors. For these results, we measured the upstream sectors based on industry- or firm-level citation networks.

At face value, our evidence implies that the bulk of the productivity slowdown in the United States (and several other industrialized economies) can be explained by the sizable increase in the cross-industry variance of TFP growth and innovation. For example, if TFP growth variance had remained at its 1977–1987 level for the subsequent two decades, the US manufacturing productivity would have grown twice as rapidly in 1997–2007 as it did—yielding a counterfactual growth rate that would have exceeded its observed level in either of the two prior decades. These estimates illustrate the potential importance of our mechanism, but given the limitations of our measurement and sources of variation, they do not constitute a definitive assessment of its quantitative contribution.

We view our paper as a first step in the theoretical and empirical investigation of the interlinked nature of innovation across sectors. Based on the findings above, many areas of research appear fruitful. 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, this initial evidence highlights the need for additional empirical strategies to explore dependencies among innovating sectors and the innovations generated by their suppliers. These same relationships could be tested, for example, using firm-level input-output data, where we suspect that the importance of supplier-customer linkages would be even larger. Third, another interesting context to explore is the role of global supply chains in productivity bottlenecks. On the one hand, imported intermediates and technologies can relax domestic bottlenecks. On the other hand, global supply chains may introduce more extensive technological dependencies, which could intensify bottlenecks if those trade channels become constrained. Fourth, it would be valuable to investigate the bottleneck hypothesis using historical data— focusing for example on major technological breakthroughs in the first half of the 20^{th} century. Finally, our framework makes a strong—perhaps even rash—prediction, whose verification awaits the passage of time: if and when lagging industries ultimately increase their innovation and productivity growth rates, a rapid takeoff in aggregate productivity should ensue.

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 Autor, David Dorn, Gordon Hanson, Gary Pisano, and Pian Shu. “Foreign competition and domestic innovation: Evidence from U.S. patents. *AER: Insights*, 2(3):357–374, 2020.

Benjamin Balsmeier, Mohamad Assaf, Tyler Chesebro, Gabe Fierro, Kevin Johnson, Scott Johnson, Guan-Cheng Li, Sonja Lück, Doug O'Reagan, Bill Yeh, Guangzheng Zang, and Lee Fleming. Machine learning and natural language processing on the patent corpus: Data, tools, and new measures. *Journal of Economics & Management Strategy*, 27(3):535–553, 2018. doi: https://doi.org/10.1111/jems.12259.

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*, 13 (1):333–72, 2021.

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.

Steven J Davis, John C Haltiwanger, Scott Schuh, et al. Job creation and destruction. *MIT Press Books*, 1, 1998.

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.

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.

Ernest Liu and Song Ma. Innovation networks and r&d allocation. Working Paper 29607, National Bureau of Economic Research, December 2021.

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

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

Table 1: 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)
Industry 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)
Industry 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: This table reports estimates of equation (6). The dependent variable is an industry's TFP growth in a five-year period and the two key right-hand side variables are mean and variance of TFP growth among that industry's suppliers. Time dummies are included in all regressions and industry dummies (corresponding to linear industry trends) are included in columns 3, 5, 7 and 8. Columns 1–5 report unweighted OLS regressions, and columns 6–8 use the industry's 1987 share of real value-added as weights. Panel A is for manufacturing industries only from 1977-2007, and Panel B is for all industries from 1987-2007. Industries are defined using 1997 NAICS codes. Standard errors are clustered at the industry level.

Table 2: Examples of Limiting and Limited Industries

Panel A: List of Select Fastest-Growing Industries that Drive Rising TFP Variance
Semiconductor and Related Devices
Electronic Computers
Iron and Steel Mills
Computer Storage Devices
Motor Vehicle Electrical and Electronic 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 (Bottlenecked) Industries
Surgical and Medical Instruments
Relay and Industrial Controls
Gasoline Engine and Engine Parts
Guided Missile and Space Vehicles
Industrial Valves

Notes: Bottleneck industries (Panel B) are defined as those for which a 10% increase in TFP would result in the *largest* aggregate reduction in the variance of TFP growth across all supplying industries (i.e. $\text{VAR}\left(\Delta TFP_{jt}\right)$ from Equation 6). Fastest-growing industries (Panel A) are conversely defined as those for which a 10% increase in TFP would result in the *smallest* aggregate reduction in the variance of TFP growth across supplying industries. Limited ("bottlenecked") industries (Panel C) are defined as the 50 manufacturing industries with the highest variance of TFP among suppliers, after limiting to the 100 industries with the highest value-added. Sample is restricted to 462 manufacturing industries during 1997–2007. See Appendix Table A2 for an ordered list of the top 10 industries in each category during 1997–2002 and 2002–2007.

Table 3: Country-Specific Instruments

	(1)	(2)	(3)	(4)	(5)	(6)
			A: Average TFP Growth			
	OLS Estimates		IV Estimates			
Upstream Average	0.951	0.780	1.369	1.416	1.387	1.509
	(0.232)	(0.119)	(0.363)	(0.655)	(0.378)	(0.758)
Upstream Variance	-0.876	-1.066	-0.902	-0.887	-0.897	-0.795
	(0.155)	(0.135)	(0.385)	(0.527)	(0.391)	(0.588)
Estimate	OLS	OLS	2SLS	2SLS	LIML	LIML
Ind. Fixed Effects	no	yes	no	yes	no	yes
Observations	2478	2478	2478	2478	2478	2478
R-Squared	0	0	0	0	0	0
First-Stage F-Stat			1.38	.63	1.38	.63
			B: Rank of TFP growth			
	OLS Estimates		IV Estimates			
Upstream Average	0.951	0.780	0.928	1.093	0.928	1.094
	(0.232)	(0.119)	(0.338)	(0.348)	(0.342)	(0.349)
Upstream Variance	-0.876	-1.066	-0.667	-1.480	-0.664	-1.482
	(0.155)	(0.135)	(0.445)	(0.661)	(0.449)	(0.665)
Estimate	OLS	OLS	2SLS	2SLS	LIML	LIML
Ind. Fixed Effects	no	yes	no	yes	no	yes
Observations	2478	2478	2478	2478	2478	2478
R-Squared	0	0	0	0	0	0
First-Stage F-Stat			.8	2.1	.8	2.1

Notes: This table reports instrumental-variables estimates of equation (6) for all industries for 1982–2007. The dependent variable is an industry's TFP growth in a five-year period and the two key right-hand side variables are mean and variance of TFP growth among that industry's suppliers. Excluded instruments are mean and variance of supplier TFP growth in France, Germany and the UK. All regressions are unweighted. Time dummies are included in all columns and industry dummies (corresponding to linear industry trends) are included in even-numbered columns. Columns 3 and 4 report two-stage least squares (2SLS) estimates, and columns 5 and 6 report limited information maximum likelihood (LIML) estimates. Panel A defines the upstream moments, taking the average and variance of TFP growth across industries. In Panel B, we rank industries in each country according to their TFP growth and calculate the input-share weighted average and variance of TFP ranks. Standard errors are clustered at the aggregated KLEMS industry level.

Table 4: Relationship between industry TFP growth and the distribution of TFP growth

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Manufacturing Industries				All Industries			
				A. Future Supplier TFP Growth				
Future Input Average	0.166	0.154	-0.006	0.083	-0.003	0.016	-0.171	-0.063
	(0.145)	(0.123)	(0.079)	(0.077)	(0.178)	(0.168)	(0.110)	(0.104)
Future Input Variance	0.065	0.011	-0.010	-0.066	0.244	0.239	0.343	0.149
	(0.098)	(0.120)	(0.101)	(0.103)	(0.133)	(0.156)	(0.158)	(0.146)
Input Average		0.787		0.670		0.867		0.752
		(0.121)		(0.076)		(0.158)		(0.116)
Input Variance		-0.810		-0.919		-0.982		-1.061
		(0.124)		(0.122)		(0.163)		(0.186)
Ind. Fixed Effects	no	no	yes	yes	no	no	yes	yes
Industry Weighting	None	None	None	None	None	None	None	None
Observations	2772	2772	2772	2772	2016	2016	2016	2016
R-Squared	0.085	0.137	0.334	0.371	0.073	0.106	0.375	0.399
				B. Customer TFP Growth				
Customer Average	0.626	0.499	0.499	0.367	0.460	0.344	0.383	0.280
	(0.066)	(0.066)	(0.074)	(0.078)	(0.077)	(0.078)	(0.102)	(0.099)
Customer Variance	-0.503	-0.302	-0.796	-0.529	-0.443	-0.264	-0.907	-0.662
	(0.200)	(0.236)	(0.161)	(0.151)	(0.257)	(0.314)	(0.265)	(0.254)
Input Average		0.466		0.454		0.765		0.687
		(0.123)		(0.083)		(0.173)		(0.122)
Input Variance		-0.566		-0.634		-0.778		-0.784
		(0.127)		(0.125)		(0.178)		(0.201)
Ind. Fixed Effects	no	no	yes	yes	no	no	yes	yes
Industry Weighting	None	None	None	None	None	None	None	None
Observations	2769	2769	2769	2769	2015	2015	2015	2015
R-Squared	0.157	0.173	0.373	0.387	0.093	0.115	0.393	0.408
				C. Lagged TFP Growth: Dependent Variable and Supplier Metrics				
Input Average	0.641	0.637	0.509	0.530	0.915	0.921	0.724	0.744
	(0.101)	(0.097)	(0.077)	(0.081)	(0.151)	(0.153)	(0.117)	(0.121)
Input Variance	-0.678	-0.715	-0.753	-0.776	-0.923	-0.939	-0.889	-1.014
	(0.111)	(0.112)	(0.121)	(0.126)	(0.152)	(0.163)	(0.173)	(0.183)
Lagged Input Average		0.056		0.208		0.069		0.330
		(0.110)		(0.085)		(0.137)		(0.112)
Lagged Input Variance		0.015		-0.471		-0.049		-0.793
		(0.146)		(0.130)		(0.195)		(0.177)
Lagged Dep. Var.	0.089	0.086	-0.255	-0.280	0.070	0.066	-0.362	-0.391
	(0.099)	(0.103)	(0.042)	(0.045)	(0.115)	(0.122)	(0.045)	(0.050)
Ind. Fixed Effects	no	no	yes	yes	no	no	yes	yes
Industry Weighting	None	None	None	None	None	None	None	None
Observations	2310	2310	2310	2310	1974	1974	1974	1974
R-Squared	0.129	0.130	0.418	0.425	0.107	0.108	0.464	0.476

Notes: This table reports estimates of equation (6). The dependent variable is an industry's TFP growth in a five-year period and the right-hand side variables are mean and variance of TFP growth among that industry's suppliers, plus lead terms, mean and variance of TFP growth among the industry's customers, and lagged dependent variables. Time dummies are included in all regressions and industry dummies (corresponding to linear industry trends) are included in columns 3, 4, 7 and 8. Columns 1–4 are for manufacturing industries for 1977–2007 and 5–8 for all industries for for 1987–2007. All regressions are unweighted. In addition to the mean and variance of TFP growth among an industry's suppliers, Panel A includes the five-year lead of the same variables. Panel B includes the mean and variance of TFP growth among the industry's customers. Panel C includes the five-year lagged mean and variance of TFP growth among the industry's suppliers and the lag of the dependent variable (the industry's TFP growth rate). Standard errors are clustered at the industry level.

Table 5: Robustness for Downstream TFP and Upstream TFP: Manufacturing Industries

	(1) Baseline	(2) Weighted	(3) 10-year	(4) Cov.	(5) China Shock	(6) No Comp.	(7) Outlier Robust	(8) Fixed IO	(9) All Inputs	(10) 3-digit Leaveout
					A: Without Industry Trends					
Input Average	0.810	0.837	0.931	0.660	0.565	0.560	0.592	0.878	1.154	0.076
	(0.130)	(0.133)	(0.182)	(0.116)	(0.123)	(0.066)	(0.043)	(0.115)	(0.163)	(0.106)
Input Variance	-0.744	-0.889	-0.477	-0.640	-0.724	-1.231	-0.774	-0.711	-0.903	-0.360
	(0.121)	(0.109)	(0.094)	(0.127)	(0.113)	(0.587)	(0.059)	(0.161)	(0.171)	(0.120)
Input Covariance				-0.069						
				(0.170)						
Industry Fixed Effects	no	no	no	no	no	no	no	no	no	no
Observations	2772	2772	1386	2310	1386	2604	2772	2772	2772	2772
R-Squared	0.133	0.168	0.100	0.123	0.126	0.122	0.208	0.153	0.147	0.080
					B: With Industry Trends					
Input Average	0.653	0.617	0.704	0.482	0.502	0.652	0.597	0.716	0.968	0.231
	(0.074)	(0.086)	(0.122)	(0.074)	(0.100)	(0.072)	(0.048)	(0.071)	(0.097)	(0.099)
Input Variance	-0.912	-0.874	-0.641	-0.647	-0.820	-1.197	-0.757	-0.961	-1.149	-0.449
	(0.118)	(0.101)	(0.107)	(0.131)	(0.146)	(0.624)	(0.068)	(0.139)	(0.163)	(0.155)
Input Covariance				-0.399						
				(0.143)						
Industry Fixed Effects	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Observations	2772	2772	1386	2310	1386	2604	2772	2772	2772	2772
R-Squared	0.371	0.434	0.549	0.385	0.472	0.252	0.487	0.379	0.378	0.338

Notes: This table reports estimates of equation (6) for manufacturing industries for 1977–2007. The dependent variable is an industry's TFP growth in a five-year period and the right-hand side variables are mean and variance of TFP growth among that industry's suppliers plus additional controls. Time dummies are included in all regressions. Panel B additionally includes industry dummies (corresponding to linear industry trends). Column 1 repeats our baseline regression from column 2 of Table 1. Column 2 weights observations by the industry's share of 1987 real value-added. Column 3 uses 10-year observations. Column 4 controls for the covariance between the supplier TFP growth in the current and the prior five-year periods. Column 5 controls for the China shock following Autor et al. (2013). Column 6 excludes the computers and electronics manufacturing sector (NAICS 334) from the regression sample and from the construction of the average and variance of TFP growth among suppliers. Column 7 runs an outlier-robust regression (`rreg`). Column 8 fixes the input-output table at 1987. Column 9 defines the input-output network to use the share among all inputs instead of among intermediaries. Column 10 excludes the industry's own three-digit NAICS code when constructing the input-output network.

Table 6: Exploring Neoclassical Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline		Prices		Employment		Combined	
Input 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)
Input 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)
Industry 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

Notes: This table reports estimates of equation (6) for manufacturing industries for 1977–2007. The dependent variable is an industry's TFP growth in a five-year period and the right-hand side variables are mean and variance of TFP growth among that industry's suppliers plus the mean and variance of supplier prices and employment. Time dummies are included in all regressions and industry dummies (corresponding to linear industry trends) are included in even-numbered columns. All regressions are unweighted. Industries are defined using 1997 NAICS codes. Standard errors are clustered at the industry level.

Table 7: Bottleneck Regressions Using Patenting by CPC Code

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Patents			DHS Specification		
Citation Average	1.291	1.335	1.473	1.284	1.307	1.402
	(0.064)	(0.064)	(0.093)	(0.065)	(0.064)	(0.098)
Citation Variance		-0.959	-0.876		-1.040	-0.911
		(0.266)	(0.314)		(0.381)	(0.532)
CPC Fixed Effects	no	no	yes	no	no	yes
Observations	4326	4326	4323	4379	4379	4376
R-Squared	0.245	0.250	0.372	0.202	0.207	0.307

Notes: This table reports estimates of equation (7) for 1975–2014. The dependent variable is log patents or the DHS transformation of patents at the CPC level during a five-year period. The two key right-hand side variables are mean and variance of log patents or DHS patents among the "idea suppliers" of that CPC, defined from the citation matrix of patents. All regressions are unweighted. Time dummies are included in all regressions and CPC fixed effects (corresponding to linear CPC trends) are included in columns 3 and 6. The Davis-Haltiwanger-Schuh (DHS) transformation is defined as $\left(\frac{P_{i,t}-P_{i,t-1}}{\frac{1}{2}(P_{i,t}+P_{i,t-1})}\right)$. Standard errors are clustered at the CPC level.

Table 8: Bottleneck Regressions using Cross-Country Variation in Patenting as Instruments for US-firm Patenting

	Panel A: Log Patents					
	(1)	(2)	(3)	(4)	(5)	(6)
Citation Average	1.335	1.473	1.495	1.788	1.497	1.802
	(0.064)	(0.093)	(0.075)	(0.144)	(0.076)	(0.148)
Citation Variance	-0.959	-0.876	-1.289	-2.289	-1.300	-2.477
	(0.266)	(0.314)	(0.452)	(1.067)	(0.469)	(1.201)
Year-by-CPC FEs	no	yes	no	yes	no	yes
Estimator	OLS	OLS	2SLS	2SLS	LIML	LIML
Observations	4326	4323	4285	4283	4285	4283
R-Squared	0.250	0.372	0.162	0.095	0.162	0.092
First Stage F-stat	0	0	14.53	8.51	14.53	8.51
	Panel B: DHS Specification					
Citation Average	1.307	1.402	1.446	1.679	1.447	1.687
	(0.064)	(0.098)	(0.076)	(0.150)	(0.076)	(0.153)
Citation Variance	-1.040	-0.911	-1.206	-2.176	-1.208	-2.288
	(0.381)	(0.532)	(0.475)	(1.097)	(0.483)	(1.179)
Year-by-CPC FEs	no	yes	no	yes	no	yes
Estimator	OLS	OLS	2SLS	2SLS	LIML	LIML
Observations	4379	4376	4325	4324	4325	4324
R-Squared	0.207	0.307	0.138	0.078	0.138	0.077
First Stage F-stat	0	0	30.8	13.73	30.8	13.73

Notes: This table reports instrumental-variables estimates of equation (7) for 1975–2014. The dependent variable is log patents (Panel A) or the DHS transformation of patents (Panel B) at the CPC level during a five-year period. The two key right-hand side variables are mean and variance of log patents or DHS patents among the "idea suppliers" of that CPC, defined from the citation matrix of patents. The excluded instruments are mean and variance of patenting among the top 10 countries with the most patents over our sample period (Canada, China, Germany, France, the UK, Italy, Japan, Korea, Russia and Taiwan). All regressions are unweighted. Time dummies are included in all regressions and CPC fixed effects (corresponding to linear CPC trends) are included in even-numbered columns. Columns 3 and 4 report two-stage least squares (2SLS) estimates, and columns 5 and 6 report limited information maximum likelihood (LIML) estimates. The Davis-Haltiwanger-Schuh (DHS) transformation is defined as $\left(\frac{P_{i,t}-P_{i,t-1}}{\frac{1}{2}(P_{i,t}+P_{i,t-1})}\right)$. Standard errors are clustered at the CPC level.

Table 9: Bottleneck Regressions using Patenting by CPC Code: Robustness to Lags, Leads, and Citing CPCs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log Patent				DHS Specification			
A. Future Patenting Growth Among Cited CPCs								
Citation Average		1.284		1.458		1.317		1.442
		(0.079)		(0.098)		(0.087)		(0.104)
Citation Variance		-0.733		-0.697		-1.100		-1.303
		(0.242)		(0.299)		(0.418)		(0.532)
Future Citation Average	0.927	0.064	0.414	0.180	0.907	0.024	0.349	0.129
	(0.067)	(0.077)	(0.089)	(0.086)	(0.073)	(0.088)	(0.113)	(0.108)
Future Citation Variance	-0.756	-0.530	-0.629	-0.351	-0.981	-0.643	-0.923	-0.647
	(0.246)	(0.233)	(0.290)	(0.278)	(0.421)	(0.505)	(0.589)	(0.568)
CPC Fixed Effects	no	no	yes	yes	no	no	yes	yes
Observations	3712	3712	3709	3709	3753	3753	3753	3753
R-Squared	0.193	0.275	0.362	0.424	0.167	0.243	0.315	0.370
B. Patenting Growth Among Citing Patents								
Citation Average		0.919		0.395		0.726		0.252
		(0.237)		(0.328)		(0.296)		(0.368)
Citation Variance		-1.003		-1.374		-1.437		-1.628
		(0.442)		(0.524)		(0.883)		(1.169)
Citing Patent Average	1.276	0.429	1.516	1.152	1.262	0.600	1.482	1.252
	(0.066)	(0.234)	(0.089)	(0.326)	(0.067)	(0.292)	(0.084)	(0.356)
Citing Patent Variance	-1.204	-0.111	-0.897	0.240	-1.159	0.296	-0.848	0.546
	(0.245)	(0.425)	(0.306)	(0.512)	(0.299)	(0.808)	(0.365)	(1.011)
CPC Fixed Effects	no	no	yes	yes	no	no	yes	yes
Observations	4326	4326	4323	4323	4378	4378	4376	4376
R-Squared	0.243	0.252	0.375	0.379	0.204	0.211	0.312	0.314
C. Lagged Patenting Growth: Dependent Variable and Citation Metrics								
Citation Average	1.278	1.260	1.460	1.394	1.365	1.225	1.526	1.428
	(0.061)	(0.081)	(0.101)	(0.102)	(0.068)	(0.092)	(0.103)	(0.105)
Citation Variance	-0.985	-0.965	-1.092	-1.017	-0.793	-0.797	-0.581	-0.488
	(0.271)	(0.269)	(0.340)	(0.341)	(0.369)	(0.399)	(0.493)	(0.491)
Lagged Citation Average		0.032		0.509		0.226		0.704
		(0.103)		(0.115)		(0.121)		(0.138)
Lagged Citation Variance		-0.061		-0.332		-0.073		-0.165
		(0.257)		(0.316)		(0.398)		(0.449)
Lagged Dep. Var.	0.055	0.052	-0.097	-0.132	-0.068	-0.084	-0.197	-0.238
	(0.034)	(0.038)	(0.033)	(0.036)	(0.042)	(0.046)	(0.040)	(0.041)
CPC Fixed Effects	no	no	yes	yes	no	no	yes	yes
Observations	3695	3695	3693	3693	3743	3743	3742	3742
R-Squared	0.256	0.256	0.399	0.406	0.224	0.226	0.377	0.388

Notes: This table reports estimates of equation (7) for 1975–2014. The dependent variable is log patents or the DHS transformation of patents at the CPC level during a five-year period. The right-hand side variables are mean and variance of log patents or DHS patents among the "idea suppliers" of that CPC, defined from its citation across CPCs, plus lead terms, mean and variance of patenting among citing (rather than cited) CPCs, and lagged dependent variables. All regressions are unweighted. Time dummies are included in all regressions and CPC fixed effects (corresponding to linear CPC trends) are included in columns 3, 4, 7 and 8. The Davis-Haltiwanger-Schuh (DHS) transformation is defined as $\left(\frac{P_{i,t}-P_{i,t-1}}{\frac{1}{2}(P_{i,t}+P_{i,t-1})}\right)$. Standard errors are clustered at the CPC level.

Table 10: Robustness for Bottleneck Regressions Using Patenting by CPC Code

	(1) Baseline	(2) Weighted	(3) 10-year	(4) Cov.	(5) No Comp.	(6) 3-digit Leaveout	(7) Excluding post-2005	(8) US-Firms Only
				A: Without CPC Trends				
Citation Average	1.335	1.305	1.378	1.330	1.377	1.175	1.292	1.299
	(0.064)	(0.084)	(0.073)	(0.065)	(0.074)	(0.078)	(0.072)	(0.060)
Citation Variance	-0.959	-1.414	-0.840	-0.730	-0.102	-0.624	-0.945	-0.323
	(0.266)	(0.393)	(0.197)	(0.298)	(0.260)	(0.219)	(0.271)	(0.218)
Citation Covariance				-0.337				
				(0.454)				
CPC Fixed Effects	no	no	no	no	no	no	no	no
Observations	4326	4305	1853	3094	3783	4326	3098	4224
R-Squared	0.250	0.442	0.347	0.279	0.265	0.195	0.230	0.207
				B: With CPC Trends				
Citation Average	1.473	1.433	1.555	1.466	1.474	1.211	1.326	1.363
	(0.093)	(0.129)	(0.112)	(0.107)	(0.087)	(0.109)	(0.113)	(0.091)
Citation Variance	-0.876	-1.215	-0.594	-0.765	-0.310	-0.574	-0.724	-0.209
	(0.314)	(0.400)	(0.247)	(0.356)	(0.318)	(0.295)	(0.370)	(0.284)
Citation Covariance				0.179				
				(0.471)				
CPC Fixed Effects	yes	yes	yes	yes	yes	yes	yes	yes
Observations	4323	4304	1846	3090	3781	4323	3096	4221
R-Squared	0.372	0.580	0.601	0.453	0.369	0.333	0.414	0.303

Notes: This table reports estimates of equation (7) for 1975–2014. The dependent variable is log patents or the DHS transformation of patents at the CPC level during a five-year period. The right-hand side variables are mean and variance of log patents or DHS patents among the "idea suppliers" of that CPC, plus additional controls. Column 1 is the baseline estimate, from column 1 of Table 8. Column 2 weights observations by the CPC code's share of total patenting in the sample period. Column 3 uses 10-year observations. Column 4 controls for the covariance between the idea-supplier patenting growth in the current and the prior five-year periods. Column 5 removes patents that belong to CPC class G, which includes computers. Column 6 excludes the observation's entire three-digit CPC code when constructing the citation network. Column 7 limits the sample to years before 2005, while column 8 limits the sample to patents from US-based firms. Time dummies are always included, and Panel B additionally includes CPC fixed effects (corresponding to linear CPC trends). The Davis-Haltiwanger-Schuh (DHS) transformation is defined as $\left(\frac{P_{i,t}-P_{i,t-1}}{\frac{1}{2}(P_{i,t}+P_{i,t-1})}\right)$. Standard errors are clustered at the CPC level.

Table 11: Bottleneck Patterns Using Firm-level Patenting

	(1)	(2)	(3)	(4)	(5)	(6)
	DHS Specification			DHS Specification with zeros		
Citation Average	1.089	1.048	1.039	0.234	0.282	0.233
	(0.010)	(0.024)	(0.015)	(0.003)	(0.006)	(0.004)
Citation Variance	-0.264	0.337	-0.292	-0.065	-0.025	-0.051
	(0.042)	(0.083)	(0.045)	(0.011)	(0.017)	(0.012)
Firm FEs	no	yes	no	no	yes	no
CPC*Year FEs	no	no	yes	no	no	yes
Observations	654583	617894	640397	1888705	1888705	1828778
R-Squared	0.037	0.414	0.044	0.009	0.038	0.013

Notes: This table reports estimates of equation (7) for 1975–2014 using firm-level observations. The dependent variable is the DHS transformation of a firm's patents during a five-year period. The two key right-hand side variables are mean and variance of patenting among the "idea suppliers" of the firm, defined from its citation across CPCs. All regressions are unweighted. Time dummies are included in all regressions. Firm fixed effects (corresponding to linear firm-level trends) are included in columns 2 and 5, and CPC-times-year fixed effects are included in columns 3 and 6. The Davis-Haltiwanger-Schuh (DHS) transformation is defined as $\left(\frac{P_{i,t}-P_{i,t-1}}{\frac{1}{2}(P_{i,t}+P_{i,t-1})}\right)$. In columns 1–3, we use the standard DHS transformation, dropping observations that have two consecutive zeros. In columns 4–6, we assign zero to the DHS when there are two consecutive zeros, thus keeping all observations. Standard errors are clustered at the firm level.

Table 12: Robustness for Bottleneck Regressions Using Firm-level Patenting

	(1) Baseline	(2) Weighted	(3) Lagged	(4) Excluding post-2005	(5) US Only	(6) Lead Horserace
			A: Without CPC-by-Year Fixed Effects			
Citation Average	1.089	2.029	0.593	1.242	1.074	1.014
	(0.010)	(0.143)	(0.011)	(0.012)	(0.012)	(0.016)
Citation Variance	-0.264	-2.567	-0.290	-0.446	-0.055	-0.187
	(0.042)	(0.407)	(0.038)	(0.057)	(0.045)	(0.059)
Lagged 5-year Growth			-0.261			
			(0.002)			
Future Citation Average						0.176
						(0.018)
Future Citation Variance						-0.341
						(0.059)
CPC*Year FEs	no	no	no	no	no	no
Observations	654583	654583	378905	384258	363911	528162
R-Squared	0.037	0.206	0.076	0.020	0.035	0.029
			B: With CPC-by-Year Fixed Effects			
Citation Average	1.039	2.391	0.591	1.209	1.018	1.064
	(0.015)	(0.133)	(0.016)	(0.019)	(0.019)	(0.022)
Citation Variance	-0.292	-2.687	-0.104	-0.462	-0.088	-0.286
	(0.045)	(0.414)	(0.041)	(0.060)	(0.049)	(0.063)
Lagged 5-year Growth			-0.253			
			(0.002)			
Future Citation Average						0.096
						(0.024)
Future Citation Variance						-0.225
						(0.063)
CPC*Year FEs	yes	yes	yes	yes	yes	yes
Observations	640397	640397	373321	380680	356228	520627
R-Squared	0.044	0.233	0.083	0.025	0.043	0.035

Notes: This table reports estimates of equation (7) for 1975–2014 using firm-level observations. The dependent variable is the DHS transformation of a firm's patents during a five-year period. The right-hand side variables are mean and variance of patenting among the "idea suppliers" of the firm, defined from its citation across CPCs, plus additional controls. All regressions are unweighted. Time dummies are included in all regressions. Firm fixed effects (corresponding to linear firm-level trends) are included in Panel B. Column 1 repeats our baseline from column 1 of Table 11. Column 2 weights each observation by the corresponding firm's share of total patenting in the sample period. Column 3 includes the lagged dependent variable on the right-hand side. Column 4 limits the sample to years before 2005. Column 5 includes only patents filed by US firms (both on the left and the right-hand side). Column 6 additionally includes the lead of the mean and variance of patenting among the idea suppliers of the firm. The Davis-Haltiwanger-Schuh (DHS) transformation is defined as $\left(\frac{P_{i,t}-P_{i,t-1}}{\frac{1}{2}(P_{i,t}+P_{i,t-1})}\right)$, and throughout this table we use the standard DHS transformation, dropping observations that have two consecutive zeros. Standard errors are clustered at the firm level.

Table 13: Evidence on Bottlenecks from Cross-Country Regressions Using TFP

	(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-by-Country FEs				X					
Year-by-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: This table reports estimates of equation (6) for 1987–2007 using cross-country observations. The dependent variable is TFP growth of an industry in a given country in a five-year period and the two key right-hand side variables are mean and variance of TFP growth among that country-industry pair's suppliers. All regressions are unweighted unless otherwise indicated. Time and country dummies are included in all regressions. The sample includes 30 industries in 9 countries: Austria, Finland, France, Germany, Italy, the Netherlands, Spain, the UK and the US. Columns 1 and 2 are cross-country analogues of columns 1 and 2 in Table 1. Column 3 includes industry-by-year fixed effects, while column 4 also adds country-by-year fixed effects. Column 5 includes the lagged dependent variable, and column 6 and 7 weight each industry observation by its share of within-country value-added (countries themselves are not weighted). Column 8 uses 10-year periods. While columns 1-8 exploit variation in input shares across both countries and industries, column 9 focuses on within-country, cross industry input-output linkages. See text for details. Standard errors are clustered at the industry level.

Table 14: Evidence on Bottlenecks from Cross-Country Regressions Using Patenting

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log Patent				DHS Specification			
Citation Average	1.231	1.240	1.229	1.404	1.103	1.096	1.077	1.200
	(0.021)	(0.022)	(0.022)	(0.054)	(0.013)	(0.013)	(0.013)	(0.033)
Citation Variance		-0.031	-0.030	-0.025		-0.048	-0.034	-0.056
		(0.010)	(0.011)	(0.018)		(0.020)	(0.020)	(0.031)
Year-by-CPC FEs	no	no	yes	yes	no	no	yes	yes
Year-by-Country FEs	no	no	no	yes	no	no	no	yes
Observations	84870	84870	84862	84862	85698	85698	85694	85694
R-Squared	0.223	0.223	0.278	0.281	0.241	0.241	0.306	0.309

Notes: This table reports estimates of equation (7) for 1975–2014 using cross-country variation at the CPC level. The dependent variable is log patents (columns 1–4) or DHS transformation of patents (columns 5–8) at the CPC level during a five-year period. The two key right-hand side variables are mean and variance of patenting among the "idea suppliers" of the country-CPC observation. We considered a 20 countries with the largest number of patents in our sample (Austria, Australia, Belgium, Canada, Switzerland, China, Germany, Denmark, Spain, Finland, France, the UK, Italy, Japan, Korea, and the Netherlands, Russia, Sweden, Taiwan, and the US). All regressions are unweighted. Time and country dummies are included in all regressions. In addition, columns 3, 4, 7 and 8 include CPC-times-year fixed effects. The Davis-Haltiwanger-Schuh (DHS) transformation is defined as $(\frac{P_{i,t}-P_{i,t-1}}{\frac{1}{2}(P_{i,t}+P_{i,t-1})})$. Standard errors are clustered at the CPC level.

9 Main Text Figures

Figure 1: Counts of US Patents Issued, 1940–2010, and Shares in (i) Electricity and Electronics and (ii) Instruments and Information

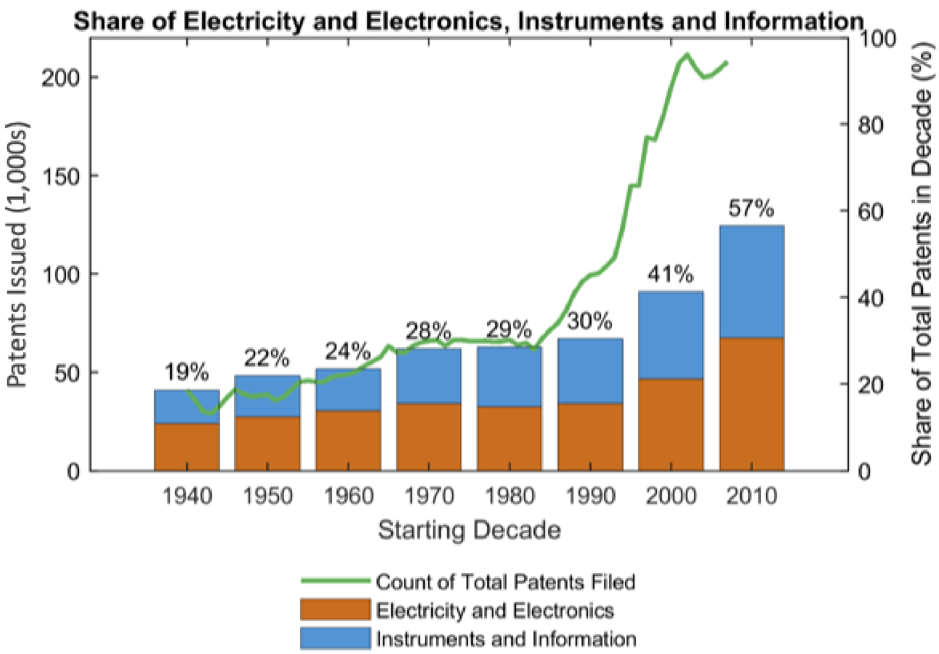

Notes: This figure plots the evolution of the counts and share (among all US utility patents) of electricity and electronics, instruments and information patents. Specifically, the left-hand y-axis (green line) gives the count of US utility patents issued in each year (green line), while the right-hand y-axis corresponds to the share of patents granted in each decade that are in Electricity and Electronics (orange bar) and Instruments and Information (blue bar). Instruments and Information is synonymous with information and communications technologies (ICT).

Figure 2: Time Series for Aggregate Total Factor Productivity

(a) United States (b) Select OECD Countries

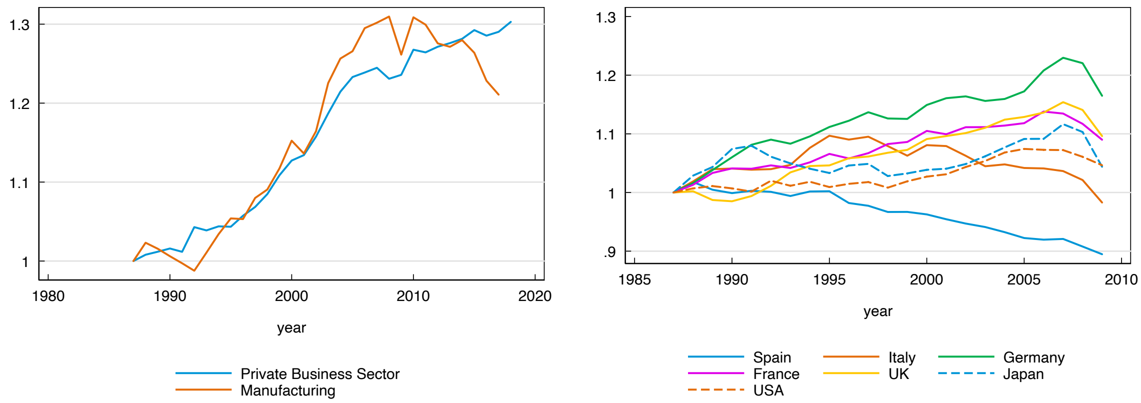

Notes: This figure plots the time series for aggregate TFP for the US private business sector and manufacturing (left panel) and for selected OECD countries (right panel). The US TFP in the left panel is normalized to 1 in 1987 and spans 1987-2017 (data from the BLS Major Sector and Major Industry Total Factor Productivity database). All TFP series are normalized to 1 in 1987 in the right panel as well and span 1987-2009 (data from the 2012 release of the EU KLEMS Growth and Productivity Accounts).

Figure 3: Bottleneck Patterns: Distribution of Upstream TFP Growth

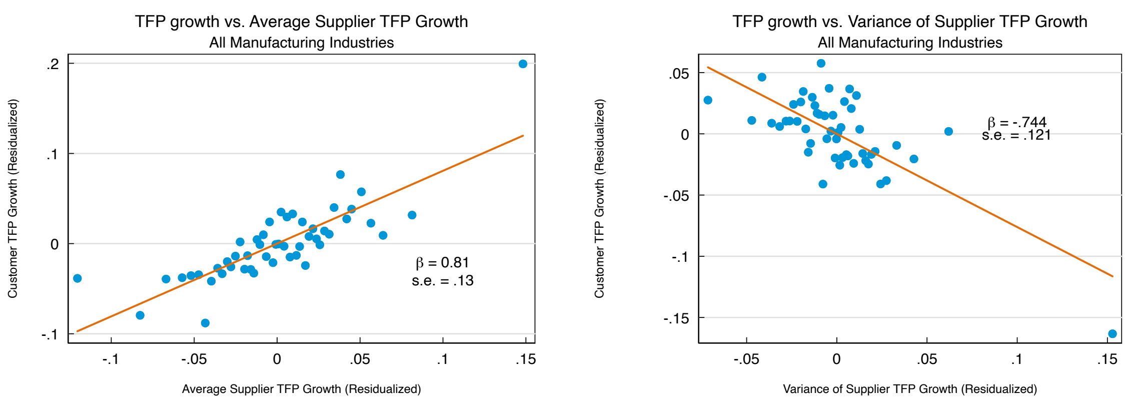

Notes: This figure reports binscatters (50 bins) for the regression model in Table 1 from Panel A, column 2 for the (conditional) relationship between manufacturing TFP growth and either the mean (left panel) or the variance (right panel) of supplier TFP growth. Specifically, the left panel plots the residuals from independent regressions of the x- and y-axis variables on the supplier *variance* of TFP growth, with time fixed effects. The right panel plots the residuals from independent regressions of the x- and y-axis variables on the supplier *average* of TFP growth, with time fixed effects.

Figure 4: Variance of TFP growth

Panel A: Variance of TFP: Manufacturing Industries

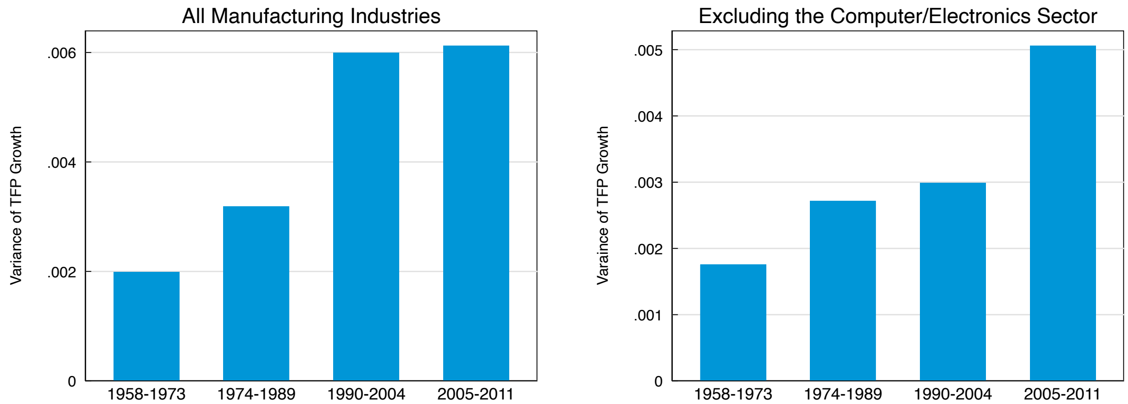

Panel B: Variance of Supplier TFP: Manufacturing Industries

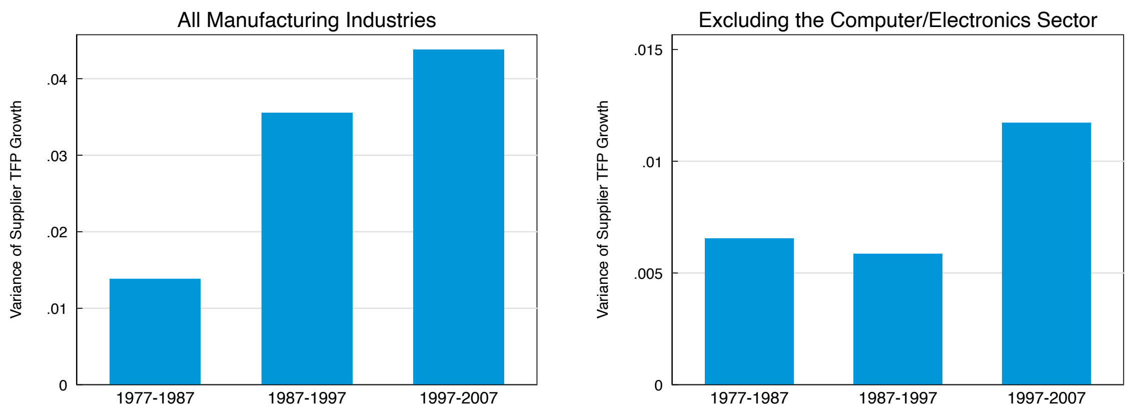

Panel C: Variance of Supplier TFP: All Industries

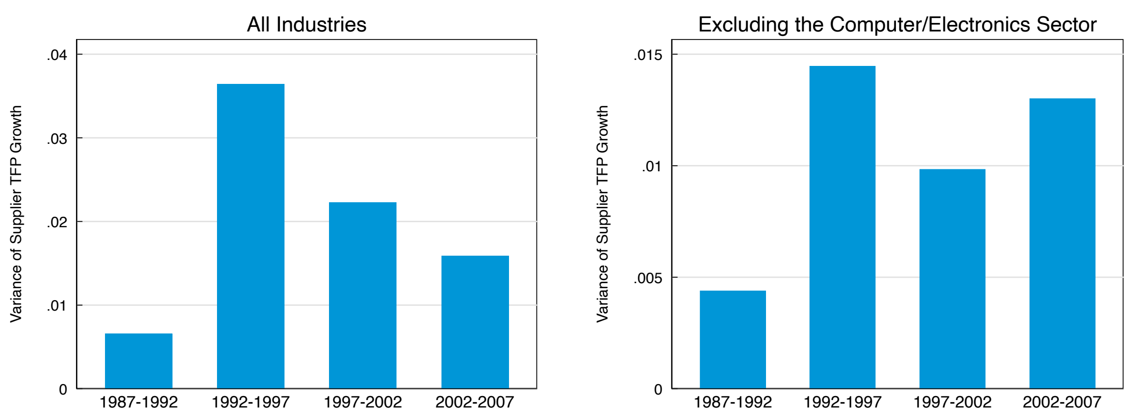

Notes: This figure plots the variance of TFP across manufacturing industries, variance of supplier TFP across manufacturing industries, and variance of supplier TFP across all industries. Each industry observation is weighted by its share of total nominal value-added. Panel A is for the variance of TFP growth across manufacturing industries for each five-year periods, spanning 1958–2011 (averaged into 15-year bars). Panel B reports the variance of supplier TFP growth across 462 six-digit, NAICS-based manufacturing industry, again for 1977–2007 (averaged into 10-year bars), while Panel C reports the variance of supplier TFP growth across all industries (adding 42 three-digit, non-manufacturing industries). Figures on the right exclude the computer and electronics sector (NAICS 334). In Panels B and C, the input-output network is defined at the beginning of each five-year period.

Figure 5: Magnitude of Bottleneck Estimates

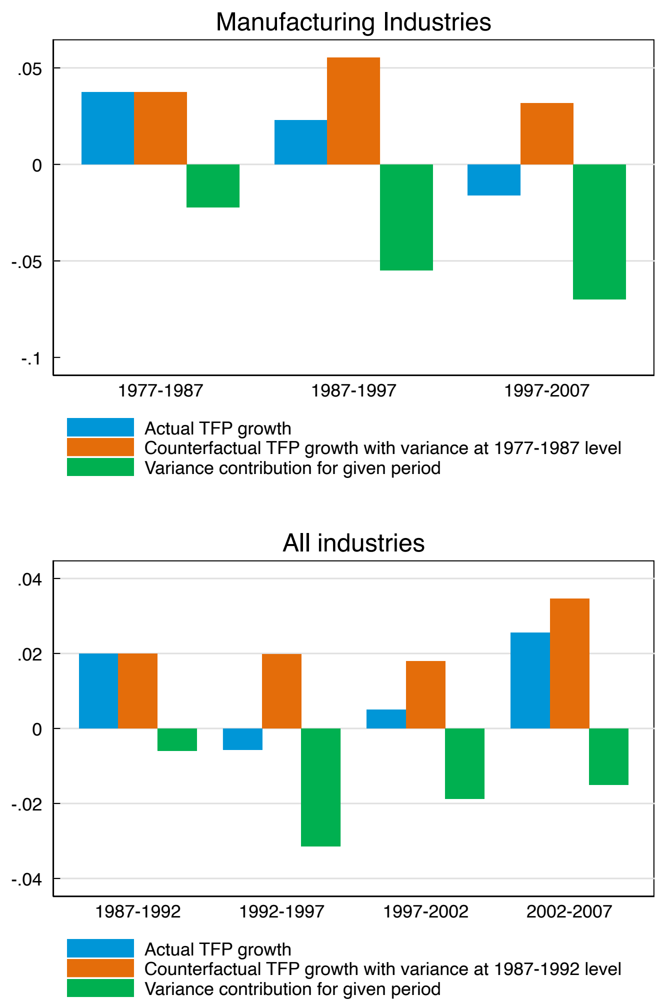

Notes: This figure reports actual and counterfactual TFP growth, and the contribution from supplier variance, for manufacturing and all industries for the periods 1977–1987, 1987–1997, and 1997–2007. The counterfactuals are based on regression estimates from the column 7 specification of Table 1. Counterfactual TFP (orange bars) is computed from the regression coefficients as the TFP growth that would have been observed in the given period if the variance of TFP growth had remained at the same level as during the initial period (1987–1992). Specifically, we calculate counterfactual TFP growth by subtracting the contribution of the increase in supplier variance relative to 1977–1987 (the 1977–1987 supplier variance is shown with the blue bar and by construction counterfactual TFP growth in 1977–1987 is equal to actual growth in this period).

Figure 6: Magnitude of Bottleneck Estimates in International Data

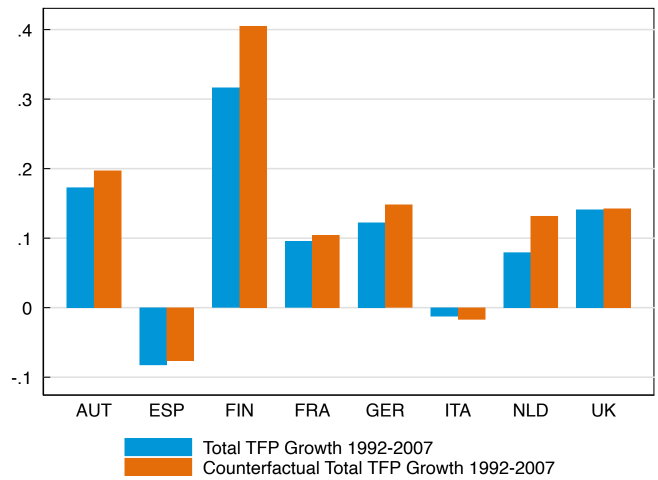

Notes: This figure reports actual and counterfactual TFP growth between 1992 and 2007 across the countries in our international panel data (Austria, Spain, Finland, France, Germany, Italy, the Netherlands and the UK). The counterfactual are based on regression estimates from the column 2 specification of Table 13. Specifically, counterfactual TFP (orange bars) is computed from the regression coefficients as the TFP growth that would have been observed in the given country and year if the variance of supplier TFP growth had remained at the same level as during the initial period (1992–1997). This is calculated by subtracting the contribution of supplier TFP variance from the actual TFP growth (blue bars).

10 Appendix Tables

Table A1: Summary Statistics Table

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

Notes: Panel A reports average growth statistics across 462 six-digit, NAICS-based manufacturing industry codes and among stacked, sequential 5-year periods during 1977–2007. Panel B reports average growth statistics across 504 industries (462 six-digit, manufacturing codes plus 42 three-digit, non-manufacturing codes) and among stacked, sequential 5-year periods during 1987–2007. Panel C reports average growth in log TFP across 30 industries in 9 countries (Austria, Finland, France, Germany, Italy, Netherlands, Spain, UK, and US) and among stacked, sequential 5-year periods during 1987–2007. Upstream metrics are calculated using intermediate cost shares from the input-output matrix. Observations (industries) are unweighted.

Table A2: Top 10 Limiting and Limited Industries

Panel A: List of Fastest-Growing Industries that Drive Rising TFP Variance	
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 Foods	Sawmills
Iron and Steel Mills	Biological Products (except Diagnostic)
All Other Plastics Products	Other Basic Inorganic Chemicals
Motor Vehicle Electrical and Electronic Equipment	Other Plastics Products
Soybean Processing	In-Vitro Diagnostic Substances
Poultry Processing	Other Basic Organic Chemicals
Motor Vehicle Metal Stamping	Petrochemicals
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	Rolled Steel Shape Manufacturing
Bare Printed Circuit Boards	Turbine and Turbine Generator Set Units
Electronic Connectors	Medicinal and Botanical Manufacturing
Other Electronic Components	Motor Vehicle Electrical and Electronic Equipment
Photographic and Photocopying Equipment	Unsupported Plastics Film and Sheets
Panel C: List of Limited (Bottlenecked) 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 Equipment	Industrial Valves

Notes: Bottleneck industries (Panel B) are defined as those for which a 10% increase in TFP would result in the *largest* aggregate reduction in the variance of TFP growth across all supplying industries (i.e. $\text{VAR}\left(\Delta TFP_{jt}\right)$ from Equation 6). Fastest-growing industries (Panel A) are conversely defined as those for which a 10% increase in TFP would result in the *smallest* aggregate reduction in the variance of TFP growth across supplying industries. Limited ("bottlenecked") industries (Panel C) are defined as the 50 manufacturing industries with the highest variance of TFP among suppliers, after limiting to the 100 industries with the highest value-added. Sample is restricted to 462 manufacturing industries during 1997–2007. See Table 2 for a select list of 5 exemplar industries from each category.

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Input 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)
Input Variance	-1.132		-1.502		-1.142		-1.379	
	(0.192)		(0.247)		(0.208)		(0.318)	
Input 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)
Industry 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: Within-industry (cross-establishment) variance comes from Dispersion Statistics on Productivity (DiSP) provided by the US Census Bureau. The sample includes stacked, sequential 5-year average changes for manufacturing industries during 1997–2002 and 2002–2007. Fixed effects for year are included in all regressions and industry fixed effects are included where indicated. For columns 1–4, observations are unweighted; for columns 5–8, each observation (industry) is weighted by its share of real value-added across all industries.

Table A4: Robustness for Downstream TFP and Upstream TFP: All Industries

	(1) Baseline	(2) Weighted	(3) 10-year	(4) Cov.	(5) China Shock	(6) No Comp.	(7) Outlier Robust	(8) Fixed IO	(9) All Inputs	(10) 3-digit Leaveout
					A: Without Industry Trends					
Input Average	0.915	0.278	1.093	0.909	0.934	0.699	0.776	0.996	0.915	0.495
	(0.161)	(0.249)	(0.233)	(0.163)	(0.168)	(0.108)	(0.075)	(0.126)	(0.161)	(0.145)
Input Variance	-0.905	-0.442	-0.638	-0.798	-1.026	-2.029	-0.984	-0.845	-0.905	-1.039
	(0.158)	(0.230)	(0.117)	(0.176)	(0.159)	(0.710)	(0.087)	(0.207)	(0.158)	(0.196)
Input Covariance				-0.228						
				(0.204)						
Industry Fixed Effects	no	no	no	no	no	no	no	no	no	no
Observations	2016	2016	1008	2016	1512	1904	2016	2016	2016	2016
R-Squared	0.102	0.022	0.070	0.103	0.116	0.103	0.170	0.119	0.102	0.083
					B: With Industry Trends					
Input Average	0.780	0.219	0.892	0.687	0.896	0.770	0.696	0.876	0.780	0.460
	(0.119)	(0.228)	(0.173)	(0.118)	(0.134)	(0.122)	(0.085)	(0.124)	(0.119)	(0.149)
Input Variance	-1.087	-0.974	-1.024	-0.907	-1.200	-1.490	-1.139	-1.139	-1.087	-1.207
	(0.191)	(0.303)	(0.251)	(0.196)	(0.214)	(0.767)	(0.104)	(0.204)	(0.191)	(0.231)
Input Covariance				-0.760						
				(0.196)						
Industry Fixed Effects	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Observations	2016	2016	1008	2016	1512	1904	2016	2016	2016	2016
R-Squared	0.399	0.365	0.665	0.406	0.471	0.291	0.589	0.409	0.399	0.384

Notes: This table reports estimates of equation (6) for all industries between 1977 and 2007. The dependent variable is an industry's TFP growth in a five-year period and the right-hand side variables are mean and variance of TFP growth among that industry's suppliers plus additional controls. Time dummies are included in all regressions. Panel B additionally includes industry dummies (corresponding to linear industry trends). Column 1 repeats our baseline regression from column 2 of Table 1. Column 2 weights observations by the industry's share of 1987 real value-added. Column 3 uses 10-year observations. Column 4 controls for the covariance between the supplier TFP growth in the current and the prior five-year periods. Column 5 controls for the China shock following Autor et al. (2013). Column 6 excludes the computers and electronics manufacturing sector (NAICS 334) both from the regression sample and from the construction of the average and variance of TFP growth among suppliers. Column 7 runs an outlier-robust regression (rreg). Column 8 fixes the input-output table at 1987. Column 9 defines the input-output network to use the share among all inputs instead of among intermediaries. Column 10 excludes the industry's own three-digit NAICS code when constructing the input-output network.

Table A5: First Stage for International Instruments for TFP Growth

	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)
Industry 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)
Industry 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: The table specifies the upstream average and variance of TFP growth statistics in France, Germany, and UK in order to instrument for downstream TFP growth in the US (as in Table 3). Panel A specifies the upstream average and variance using TFP growth. Panel B ranks industries in each country according to their TFP growth and specifies the upstream average and variance using the input-share weighted mean and variance of TFP ranks (then 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 A6: Evidence on Bottlenecks from Cross-Country Regressions Using Domestic and Foreign TFP

	(1)	(2)	(3)	(4)
Upstream Domestic Average	0.144	0.148	0.127	-0.130
	(0.081)	(0.086)	(0.046)	(0.111)
Upstream Domestic Variance		-0.523	-0.417	-0.424
		(0.165)	(0.101)	(0.107)
Upstream Foreign Average	0.178	0.151	0.347	0.396
	(0.114)	(0.154)	(0.205)	(0.217)
Upstream Foreign Variance		-0.025	0.226	0.355
		(0.186)	(0.234)	(0.277)
Year FEs	yes	yes	no	no
Country FEs	yes	yes	yes	no
Year-by-country FEs	no	no	no	yes
Year-by-Industry FEs	no	no	yes	yes
Observations	3647	3647	3647	3647
R-Squared	0.066	0.073	0.361	0.393

Notes: Standard errors are clustered at the industry level. All regressions include stacked, sequential 5-year changes during 1987–2007 for 30 industries and 9 countries (Austria, Finland, France, Germany, Italy, Netherlands, Spain, US, and UK). Upstream domestic average and variance of TFP growth are calculated across industries (within-country), weighting by the input share of each industry in 2000. Upstream foreign average and variance are calculated across industries for all other countries in the sample, weighting by the input share of the country-industry pair.

Table A7: Bottleneck Results for United States, Aggregating to KLEMS codes

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	yes	yes	yes	yes	yes	yes	yes	yes
Industry FEs	no	no	no	no	yes	yes	yes	yes
Industry Weighting	None	None	Nom. VA	Nom. VA	None	None	Nom. VA	Nom VA
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. The input-output table is defined using the GGDC World Input-Output Database for the year 2000, limited to only the US. All regressions include stacked, sequential 5-year changes during 1987–2007 for the 30 KLEMS industry classifications used for the cross-country regressions reported in Table 13. Each observation (industry) is weighted by its share of nominal value-added, across all industries, in 1987.