TNIT News

TOULOUSE NETWORK FOR INFORMATION TECHNOLOGY

INNOVATION

The search for new ideas

Issue 20 June 2019

Toulouse School of Economics | Digital Center

Microsoft

Innovate or die?

Jacques Crémer

For many observers, the speed and dynamism of 21st-century high-tech economies offer ample evidence in support of Joseph Schumpeter's belief that innovation is essential to economic development. But how can policymakers encourage innovation? And who is likely to benefit? In this issue of TNIT News, we feature the work of two TNIT researchers whose empirical analysis focuses on the causes and consequences of innovation.

Harvard Business School's **Josh Lerner** tests the view that the Chinese economic miracle undermines the idea that effective institutions, such as intellectual property rights, are essential for innovation. MIT's **Heidi Williams** investigates how US patent allowances affect firm performance and worker pay. Beyond simply raising average earnings, she finds that patents exacerbate within-firm inequality.

In two related articles, Stanford's **Nicholas Bloom** suggests that the age of innovation may have already peaked. Today the developed world has more researchers than ever, but ideas productivity - the number of ideas being produced per researcher - has been falling for decades. At the same time, the meteoric growth of cloud computing over the past 10 years appears to be slowing. Despite rising spending on cloud and R&D, new customers and new ideas appear to be in short supply.

Innovation, intellectual property and China

by Josh Lerner
(Harvard)

Does the Chinese economic miracle undermine the idea that institutions are essential to protect innovation and encourage growth? Josh Lerner is the Jacob H. Schiff Professor of Investment Banking at Harvard Business School and a leading expert on innovation policies and their impact on firm strategies. Here, he presents his recent research on intellectual property rights and innovation in China.

Harkening back to Schumpeter, the literature on law and finance argues that effective legal and financial institutions lead to better economic outcomes (see panel). But China's rapid economic growth in the past 30 years raises questions about this view. As Allan, Qian, and Qian (2005) point out, China has achieved the fastest sustained growth in history despite having poor legal and financial institutions. Instead, the economic development model China has followed in the past 30 years relies on a strong state sector, with many state-owned enterprises and strong government direction. This experience suggests that, at least in the case of China, good institutions may not be necessary for economic development, and poses profound questions for academics and policymakers alike.

The Chinese puzzle

In an ongoing line of research, Lily Fang, Chaopeng Wu, and myself examine these big questions by focusing on innovation, an activity that Schumpeter identified as critical to economic change. In this summary, I will highlight our work on intellectual property right (IPR) protection and innovation in China. In a paper co-authored with Lily Fang and Chaopeng Wu, *'Intellectual property rights protection, ownership, and innovation: Evidence from China' (2017, Review of Financial Studies)*, we shed light on the following three questions:

1. *Where has China's innovation taken place: in state-owned enterprises (SOEs) or in private-sector firms?*
2. *Are legal institutions - in particular, IPR protection - important for innovation within China?*
3. *If so, are SOEs or private sector firms more sensitive to IPR protection?*

China's ability to innovate is not only an interesting and relevant question for economists, but also a timely matter of first-order policy importance to the Chinese. Since China's economic reforms started in the late 1970s, the country's growth has largely relied on cheap labor and state-led investments in physical infrastructure. But as China's labor costs have surged and growth rates have declined in recent years, this growth model has been widely seen as obsolete. China's top leaders are promoting innovation as the key to the country's sustained economic growth: for instance, in the 13^{th} Five-Year Plan released in March 2016, innovation was listed as the first guiding principle of economic policy. But the extent to which the state can drive innovation without sound institutions and economic incentives remains in question.

Was Schumpeter right?

There are two competing hypotheses about where innovation is likely to take place and the importance of IPR protection in China. Entrepreneurs' incentives to innovate - what Schumpeter terms the "entrepreneurial spirit" - depend on their ability to capture the profit from innovation, which in turn depends on IPR protection

> **We find that innovation increases significantly after Chinese firms are privatized. On average, firms' patent stock increases by 200% to 300% in the five years after privatization compared to the five years before. The increase in innovation is significantly larger in cities with high IPR protection than in cities with low IPR protection**

and institutions such as the patent system. The danger of ex post expropriation as a result of poor IPR protection will deter innovation, consistent with classic arguments by Ken Arrow *('Economic welfare and the allocation of resources for invention' in The Rate and Discretion of Inventive Activity, edited by R. Nelson, 1962)*. This line of reasoning concludes that, in China, precisely because private firms face a high risk of expropriation, institutional quality such as IPR protection standards should be particularly important for innovation in the private sector. We call this the "Schumpeterian view".

On the other hand, despite the country's poor record of IPR protection, China has in recent years become the most prolific patent-filing country in the world *(see Figure 1)*. One explanation for this paradox is that in the absence of legal protection, state ownership acts as an alternative mechanism, giving state-owned firms both incentives (or directives) for innovation and protection against expropriation. SOEs in China, in fact, have a two-tiered defense against expropriation: through administrative measures by the government (the firms' owners), and through the courts, which are often biased in their favor. This explanation suggests that China's innovation should be led by the SOEs, and because they rely on the state, institutions such as IPR protection do not matter much. We call this the "alternative mechanisms view".

Empirical challenges

To test these hypotheses, we compare the firm-level innovation (based on patent activity) of SOEs and private firms across Chinese cities with varying levels of IPR protection. We do not, however, undertake a simple cross-sectional comparison, because doing so raises two concerns. First, SOEs and private firms are inherently different: their geographic and industry distribution is non-random and may be related to the quality of local IPR protection. Second, even the quality of local IPR protection itself can be a consequence of local innovative activities (and hence demand for IPR protection), rather than having a causal effect on innovation.

To address these empirical challenges, we exploit China's privatizations of SOEs. The idea is that the privatization events result in a sharp change in the firms' ownership structures and state affiliations, while keeping other firm attributes fixed. We can therefore compare the rates of innovation before and after the change in ownership within the same firm. By studying this before-and-after difference in innovation rates across firms in regions with varying local IPR protection standards, we can identify the joint effect of ownership type and IPR protection. In essence, these events allow us to use a difference-in-difference method.

Our approach would be problematic, however, if innovative firms and entrepreneurs felt shackled by state ownership and initiated privatizations precisely in order to engage in more innovation.

Fortunately, this concern is allayed by China's political economic history. SOE privatizations and restructurings were key policy initiatives of China's top leaders from 1996 to 2005 (the 10-year period covering the 9th and 10th Five-Year Plans). This policy drive led to a massive and sweeping privatization wave, which by some estimates ultimately privatized two-thirds of the state sector. The overarching goal of these privatizations was to increase the efficiency of China's vast state sector and to transition the country from central planning to a market orientation. In contrast, innovation became a policy focus quite recently.

> **We also find evidence that patents of private-sector firms are cited more often and have a greater international presence - in other words, are of higher quality - than patents of state-owned enterprises, suggesting that the increase in patent filings is not a consequence of 'window dressing'**

Key results

We document three main findings. First, innovation increases significantly after firms are privatized. On average, firms' patent stock increases by 200% to 300% in the five years after privatization compared to the five years before. Second, the increase in innovation is significantly larger in cities with high IPR protection than in cities with low IPR protection. A one standard-deviation increase in local IPR protection score nearly quadruples the post-privatization increase in patent stock. Third, we find evidence that patents of private-sector firms are cited more often and have a greater international presence - in other words, are of higher quality - than patents of SOEs, suggesting that the increase in patent filings is not a consequence of "window dressing". In sum, our evidence is strongly supportive of the Schumpeterian view that institutions matter, even in China. It is inconsistent with the alternative mechanisms view delineated above.

In short, China's spectacular growth poses many puzzles to the economists. We hope that this work - as well as our follow-on work on subsidies for innovation in China (see panel) - will shed light on some of these.

ENTER THE DRAGON

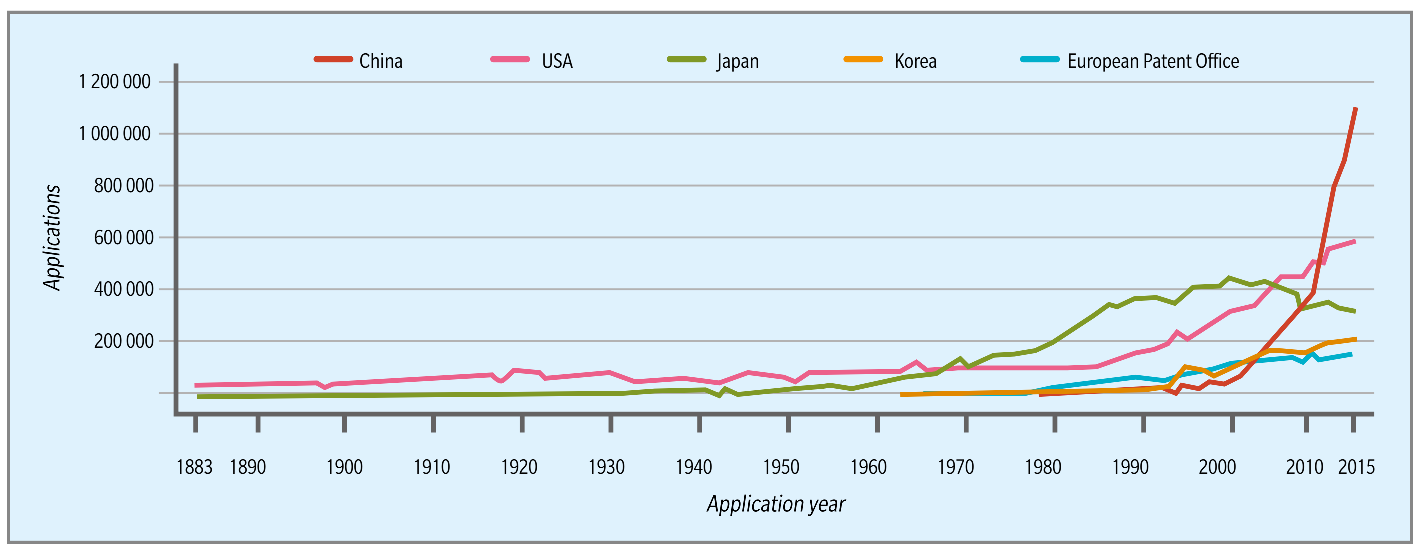

Figure 1: *rent in patent applications for the world's top five patent offices. Despite its poor record of IPR protection, China has in recent years become the most prolific patent-filing country in the world* - *Source: World Intellectual Property Organization*

THE POWER OF INSTITUTIONS

For more on the economic impact of legal and financial institutions, see *'Finance and growth: Schumpeter might be right' by R. King, and R. Levine (1993, Quarterly Journal of Economics); 'Law and finance' by R. La Porta, F. Lopez-de-Silanes, A. Shleifer, and R. Vishny (1998, Journal of Political Economy); and 'Financial dependence and growth' R. Rajan, R. and L. Zingales (1998, American Economic Review).*

SUBSIDIES AND INNOVATION

In a 2018 follow-up study, *'Corruption, government subsidies, and innovation: Evidence from China' (NBER Working Paper No. 25098)*, Josh Lerner and his coauthors continue their empirical efforts to piece together the Chinese puzzle. *"Governments are important financiers of private-sector innovation,"* they write. *"While these public funds can ease capital constraints and information asymmetries, they can also introduce political distortions. In China, where a quarter of firms' R&D expenditures come from government subsidies, we find that the anticorruption campaign that began in 2012 and the departures of local government officials responsible for innovation programs strengthened the relationship between firms' historical innovative efficiency and subsequent subsidy awards and depressed the influence of their corruption-related expenditures. We also examine the impact of these changes: subsidies became significantly positively associated with future innovation after the anti-corruption campaign and the departure of government innovation officials."*

Who profits from patents?

by Heidi Williams (MIT)

Recently picked by The Economist as one of the decade's eight best young economists, Heidi Williams is particularly interested in the causes and consequences of innovation. Here, she presents a recent paper in which she investigates how patent allowances affect firm performance and worker pay.

In standard competitive models of the labor market, we think of firms as being price takers. That is, workers are paid a wage that is a function of their skill level, and firms take market-level prices of skill as given. However, there is growing empirical evidence that firms contribute substantially to wage inequality across identically skilled workers. Put simply, how much you earn seems to depend in part on the firm at which you work (as opposed to depending solely on your skills).

One natural explanation is that perhaps firm performance matters for worker pay, in the sense that workers employed at firms that are doing better might earn more. However, testing for causal evidence on whether firm performance matters for worker pay has been challenging for two reasons.

First, from an empirical perspective, we would ideally isolate clear shocks to firm performance, and trace through how those shocks propagate into worker pay. Although that thought experiment is simple enough to describe, nearly all past attempts to analyze this question have instead analyzed observed fluctuations in firm performance over time, without making an attempt to understand their source.

Second, in the small number of cases where researchers have successfully identified clear shocks researchers have lacked the type of detailed data needed to cleanly analyze wage responses among incumbent workers. That type of data is critical because the composition of workers employed at a firm may change in response to firm-level shocks. To be concrete, when a firm discovers a new invention it may hire more skilled workers to develop and market that invention. Average wages at the firm would then go up, but that could solely reflect a compositional change in the average skill level of workers employed at the firm, even if no rents from the innovation were being passed through to worker wages.

This piece summarizes a recent academic paper that I wrote with Patrick Kline (UC-Berkeley), Neviana Petkova (US Treasury) and Owen Zidar (Princeton), in which we investigate how patent allowances affect firm performance and worker pay using a new linkage of US patent applications to US Treasury business and worker tax records.

Empirical approach: Comparing accepted and rejected patent applications

Patents provide firms with a temporary period of market power, during which they can charge supra-competitive prices and earn rents which allow them to recoup the fixed costs of their research investments. Our idea in this paper is to try to isolate quasi-random variation in which firms receive patents, and to leverage that variation in order to look at how patent-induced rents propagate into worker wages.

Specifically, consider the following thought experiment. Take two patent applications submitted by two separate firms to the US Patent and Trademark Office (USPTO) in the same year, covering the same general type of technology (in USPTO parlance, the two patent applications are sufficiently similar that they will be reviewed in the same Art Unit, or specialized group of USPTO examiners). One of the two applications is initially allowed (that is, granted on the first round of review) whereas the second application is initially rejected. We can, under some assumptions, use the initially rejected firm as a comparison for what would have happened to firm and worker outcomes at the initially allowed firm in the absence of the patent being granted.

Of course, a priori, it isn't clear that this thought experiment offers a clean comparison: it may be that better patent applications are more likely to be granted patents, in which case initially rejected firms might not be a good comparison for initially accepted firms. My research with Bhaven Sampat *('How Do Patents Affect Follow-on Innovation? Evidence from the Human Genome', American Economic Review, 2019)* has documented a potentially large idiosyncratic component to patent grants - namely, variation across patent examiners in their likelihood of granting patents to observably similar applications. More directly relevant to our study, we can assess empirically whether initially accepted and initially rejected firms look similar in terms of the levels and trends in their outcomes in the years prior to patent applications being submitted, and we find that they do, lending credibility to this empirical approach.

Data

Our empirical analysis relies on a new linkage of two datasets: the census of published patent applications submitted to the USPTO between roughly 2001-2011, and the universe of US Treasury business tax filings and worker earnings histories drawn from W2 and 1099 tax filings.

In the US, we are able to observe both accepted and "rejected" patent applications filed since 29 November 2000 under the American Inventors Protection Act. This data is what enables us to analyze the thought experiment described above, because we observe all firms filing applications, including those firms granted patents as well as those firms not granted patents. In practice, constructing this data on US patent application filings is complicated, as it requires stitching together several different USPTO administrative datasets. But in the end, we are able to combine several different public-use files from the USPTO to construct a comprehensive dataset on applications filed over this period, including information on the timing and content of the USPTO's initial decisions on each application, which is what we need to implement our empirical analysis.

We link the firms applying for patents (the so-called patent assignees) with firm names in the US Treasury business tax filings (form 1120 for C corporations, 1120S for S corporations, and form 1065 for partnerships). The business tax filings data offer a high-quality set of firm-level variables, from which we are able to construct multiple measures of firm performance. We then link these business tax filings with worker-level W2 and 1099 filings to measure the number of employees and independent contractors, as well as various worker compensation measures. The combination of the business and worker tax filings provide a window into compensation outcomes for many different types of workers, including firm officers and owners, who prevail at the top of the income distribution.

Identifying valuable patents

It is well known that most patents generate little ex post value to the firm. In our context, this means that considering the full universe of patent grants would provide very little insight into the relationship between firm-level outcomes and worker-level earnings, as most patent grants generate no shifts in firm-level outcomes. With that concern in mind, we designed our analysis to focus - in two ways - on a subsample of valuable patents which we expect, ex ante, to induce meaningful shifts in firm outcomes at the time they are allowed.

First, following the work of Farre-Mensa, Hegde, and Ljungqvist *("What Is a Patent Worth? Evidence from the US Patent 'Lottery'," 2017, NBER working paper no. 23268)*, we restrict our analysis to firms applying for a patent for the first time, for which patent decisions are likely to be most consequential.

Second, among this sample of first-time applicants, we build on the analysis of Kogan et al. *('Technological Innovation, Resource Allocation, and Growth', Quarterly Journal of Economics, 2017)* to identify a subsample of ex ante valuable patents. Kogan's team use event studies to estimate the excess stock-market return realized on the grant date of US patents assigned to publicly traded firms. We develop a methodology for extrapolating Kogan's patent value estimates to both the non-publicly traded firms in our sample and the firms whose patent applications are never granted. Specifically, we use characteristics of firms and their patent applications that are fixed at the time of application as the basis for extrapolating patent values.

Figure 1 documents that these predicted value estimates are strong predictors of treatment effect heterogeneity in our sample. Each point in ***Figure 1*** quantifies our treatment effect (using the accepted/rejected variation described above) of patents of a given value (as

PREDICTING VALUE

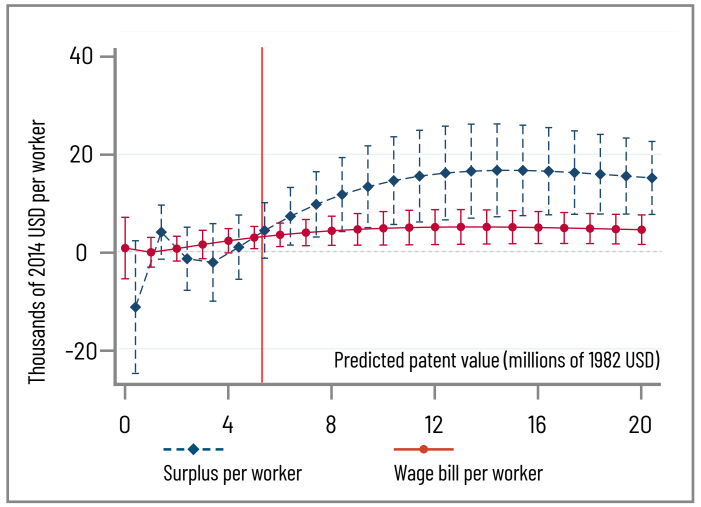

Figure 1: *Impacts by Predicted Patent Value: Surplus and Wage Bill (Re-printed version of Figure 2 from Kline et al. 2018)*

measured on the x-axis) on two different outcomes on the y-axis: surplus per worker (one measure of firm performance) and wage bill per worker (one measure of worker compensation).

As expected, and comfortingly from the perspective of validating our empirical approach, low-value patents induce essentially no changes in either firm or worker outcomes. In contrast, patents with ex ante predicted values above roughly the 80th percentile of the predicted value distribution (denoted by the red vertical line in *Figure 1*) have larger, statistically significant treatment effects on both firm and worker outcomes. Given the pattern observed in *Figure 1*, our empirical analysis pools the bottom four quintiles together and focuses on estimating the impacts of patents in the top quintile of ex ante predicted patent value.

Figures 2 and 3 document event study estimates for these same two outcome variables: surplus per worker ***(Figure 2)*** and wage bill per worker ***(Figure 3).*** Comfortingly from the perspective of validating our empirical approach, firms whose patent applications are initially allowed exhibit similar trends in firm and worker outcomes to firms whose patent applications are initially rejected in years prior to the initial decision. However, surplus per worker rises differentially for firms whose high-value patent applications are initially allowed after the initial decision date, and remains elevated afterwards ***(Figure 2)***. Similar, although more muted, trends are observed for wage bill per worker ***(Figure 3)***. The ratio of these two impacts is roughly one-third. That ratio can be interpreted as implying that workers capture roughly 30 cents of every dollar of patent-induced surplus in the form of higher earnings.

Our paper also documents evidence that, beyond simply raising average earnings at these firms, patents exacerbate within-firm inequality on a variety of margins. We find that earnings impacts are concentrated among employees in the top quartile of the within-firm earnings distribution and among employees listed on firm tax returns as "firm officers". Similarly, the earnings of owner-operators rise more than the earnings of other employees. Earnings of male employees rise strongly in response to a patent allowance, while earnings of female employees are less responsive. Inventor earnings (defined as the earnings of employees ever listed as inventors on a patent application) are more responsive than are the earnings of non-inventors, although we find substantial wage changes even for non-inventors.

Take-aways

Our paper interprets this set of empirical findings in the context of a simple economic model in which incumbent workers - that is, workers who were present at the firm at the time the patent application was filed - are imperfectly substitutable with new hires. We think this economic model appropriately captures the features of the small, innovative firms we study: the innovative work conducted at these firms is necessarily specialized and proprietary in nature, likely making it costly to replace incumbent employees with new hires. In this model, firms choose to share economic rents with incumbent workers to increase the odds of retaining them. We document empirical results consistent with that prediction: namely, worker retention rises most strongly among groups of workers with the largest earnings increases.

More broadly, our empirical findings provide some of the first evidence that truly idiosyncratic variability in firm performance is an important causal determinant of worker pay. Given that firm productivity is highly variable and persistent, it is plausible that firm-specific shocks contribute substantially to permanent earnings inequality across identically skilled workers.

PATENT POWER

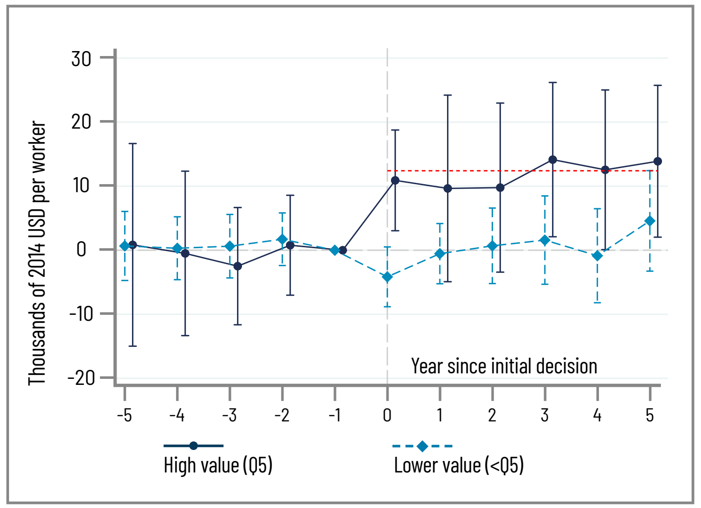

Figure 2: *Event Study Estimates: Surplus (EBITD+Wage Bill) per Worker (Re-printed version of Figure 3 from Kline et al. 2018)*

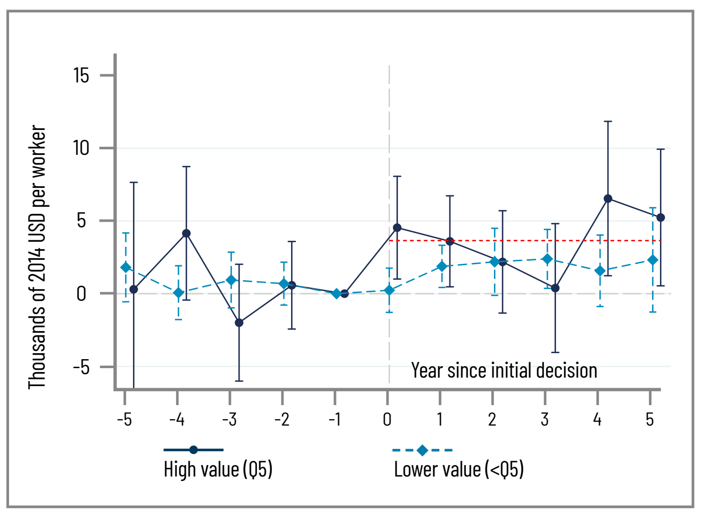

Figure 3: *Event Study Estimates: Wage Bill per Worker (Re-printed version of Figure 4 from Kline et al. 2018)*

Can cloud computing keep growing?

by Nick Bloom

(Stanford)

The lightning growth of cloud computing may have begun to wane, but for its proponents the end of the 'decade of cloud' has a silver lining. New research by Stanford economist Nicholas Bloom shows that while the uptake of cloud by new firms has slowed, companies are ramping up their spending on the technology. This corresponds to a classic theory of the diffusion of innovation, in which a new technology first spreads across firms or consumers before deepening in its use.

Cloud computing as a concept was born in the mid-2000s, with the first recorded mention of "cloud computing" arising apparently in 2006 (see Regalado, "Who coined cloud computing?", MIT Technology Review, 2011). Since that time cloud computing has seen an incredible rate of adoption across the US economy. Unlike many previous technologies such as the PC and the internet, the early adopters of cloud have been firms that are both small and large, young and old, earning cloud a moniker as a "democratic technology" (see my 2018 TNIT white paper, coauthored with Nico Pierri, "Cloud computing in the US: Democratizing Innovation").

Chasing the cloud

To evaluate the spread of cloud, I decided to look at three sources. First, adoption rates in US firms, as shown in Figure 1. The dataset for this analysis comes from Aberdeen Information's call center, which has been making annual phone calls to millions of firms across the US since the 1980s. They painstakingly record the hardware and software used by millions of firms per year. This dataset is often used by academic researchers because of its broad coverage and high quality[1]. The records from more than 150,000 US firms clearly show a massive uptake in cloud computing since 2010, but also an apparent flattening off in growth of uptake since about 2016. *Figure 2* shows this uptake by broad industry category and we can see a similar pattern.

Second, we analyzed mentions of cloud in the media, using Access World News which contains around 2,500 daily US newspapers going back to 2000. We could not directly search for the frequency of the word "cloud", since this ended up tagging weather articles

NEW HORIZONS

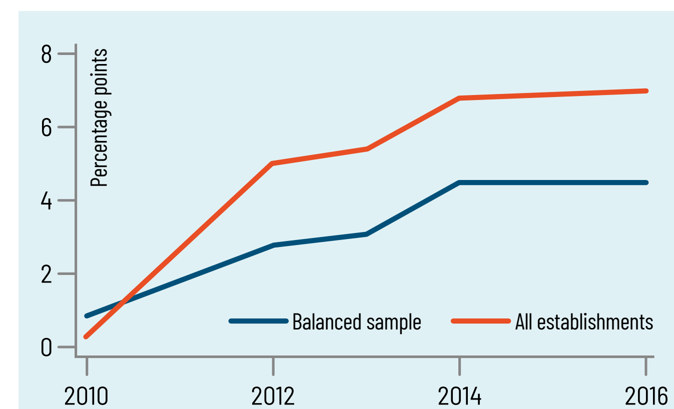

Figure 1: *Cloud Computing usage in the US has risen 20 fold since 2010, but by 2016 growth is flattening off.*

(1) This database, formerly known as Harte Hanks (or Computer Intelligence InfoCorp) has been used in small samples in a number of prior papers – for example, Bresnahan, Brynjolfsson and Hitt (2002) and Bloom, Draca and Van Reenen (2015).

> Unlike many previous technologies such as the PC and the internet, the early adopters of cloud have been firms that are both small and large, young and old, earning cloud a moniker as a "democratic technology"

USING THE CLOUD

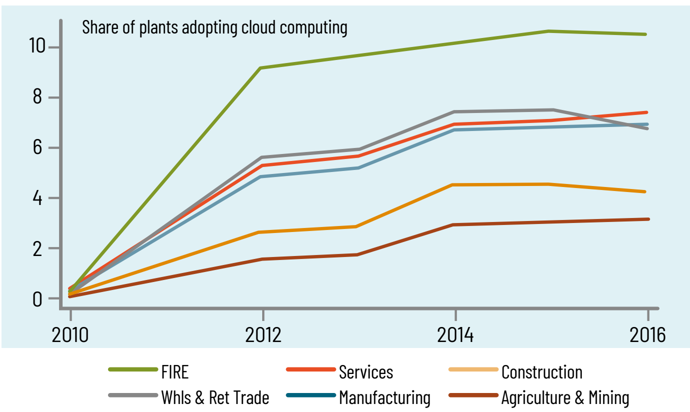

***Figure 2:** Cloud Computing usage growth appears to be flattening across all major US industries*

and uses like "His change in opinions clouded my view of the future…".

So our approach, following on from Hasan et al. ('Firm-Level Political Risk: Measurement and Effects', 2019), was to search for two-word combinations as these are typically much more exclusive in meaning and therefore less likely to generate false positives. After extensive analysis we focused on around 30 cloud-computing bigrams, with the 15 most popular ones shown in Figure 3. These included combinations like "cloud based", "cloud computing", "cloud services", "software as a service", "cloud services" and "cloud solutions".

Figure 4 pulls the frequency of cloud-computing bigrams in the press into a time series, showing a striking trend in terms of the percentage of articles mentioning a least one cloud bigram, rising from zero before 2006 up to about 0.6% by 2016. This striking rise shows how commonly cloud computing now seems to be discussed in the media, appearing as a term in about 1 in every 200 media articles. However, it also highlights how the growth in media coverage of cloud computing appears to be flattening off. According to media coverage, cloud was a new technology of the late 2000s and early 2010s. So by 2016 cloud computing had arrived - much like the cell-phone, internet or the PC - and it now appears to be viewed as a mainstream, widespread technology.

Finally, we also scraped the complete set of all company earnings

TECH TALK

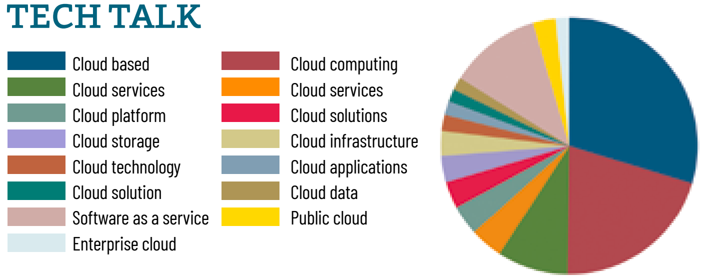

***Figure 3:** Most common bigrams (two words) for cloud computing in the media 2006-2019 inclusive.*

call reports - the approximately 90-minute earnings calls that companies take with investors every quarter when they release their latest accounting earnings data. This text is useful for understanding what topics come up in a semi-structured business conversation. Conversations are defined as "semi-structured" because while the companies clearly are trained in advance of these calls and have rehearsed responses to many questions, the analysts are free to ask whatever they like and calls can sometimes veer off into unexpected topics. In Figure 5 we show the frequency of earnings calls that mention any of our bigrams used in Figures 2 and 3 (words like "cloud computing" and "software as a service"). We can see a rapid rise in business discussions of cloud computing from around 2008 to 2013, the period of the "cloud explosion" when this highly unusual new technology quickly became commonplace. As before, the earnings calls suggest that by 2016 the adoption rate of cloud technology began to slow down.

The curse of success

To sum up, cloud computing went from being one of the hottest new technologies around 2006 to widespread use and slowing rates of adoption by 2016. We should call this the "decade of cloud". Does this mean the growth of cloud technology is over? No, far from it. Although the rates of new adoption of cloud appear to be slowing down, there is still rapid deepening of the use of the technology. (Clearly, not all cloud usage is the same. For example, firms could use cloud as a backup to cover 10% of their storage and compute costs, or as their primary application with on-premise compute as the 10% backup.) Gartner figures show exactly this is happening: Table 1 below reports cloud expenditure as a share of total IT expenditure, revealing that cloud is quickly becoming a central technology.

As a result, total expenditure on cloud is accelerating, as shown in Table 2. This S-shaped pattern should be expected for any successful technology. The "S-shaped diffusion curve" was first popularized by Everett Rogers in his 1962 book Diffusion of Innovation and shows how growth rates typically start very low with initial enthusiastic adopters, pick up rapidly as the technology spreads to the mass market, then slow down towards the end as only late adopters are left converting. Cloud, like any other new technology, has been moving through the S-curve in its first decade, and appears to now be reaching sufficiently high levels of market saturation that future adoption rates are flattening off, with growth now coming from firms deepening their use.

This is in many ways the curse of success. Cloud has become so seductive and popular - and its uptake so rapid - that the industry is now running out of new customers, so the focus appears to be on the improvement of the technology and a deepening of its use.

MUSHROOM CLOUD: AFTER THE EXPLOSION

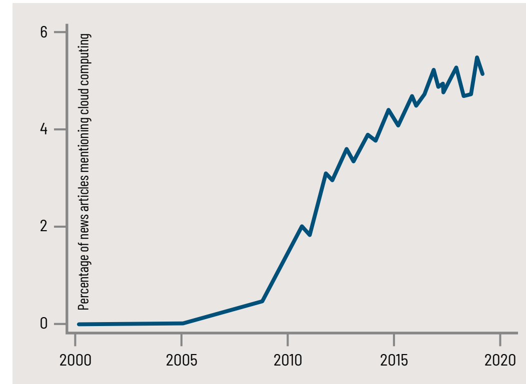

***Figure 4:** Cloud computing coverage in the media exploded after 2006, but has been flattening since 2016*

	2018	2019	2020	2021	2022
System infrastructure	11%	13%	16%	19%	22%
Infrastructure software	13%	15%	17%	18%	20%
Application software	34%	36%	38%	39%	40%
Business process outsourcing	27%	28%	29%	29%	30%
TOTAL	**19%**	**21%**	**24%**	**26%**	**28%**

Source: Gartner (August 2018) [2]

***Table 1:** Cloud shift proportion by category*

> The 'S-shaped diffusion curve' shows how growth rates typically start very low with initial enthusiastic adopters, pick up rapidly as the technology spreads to the mass market, then slow down towards the end as only late adopters are left

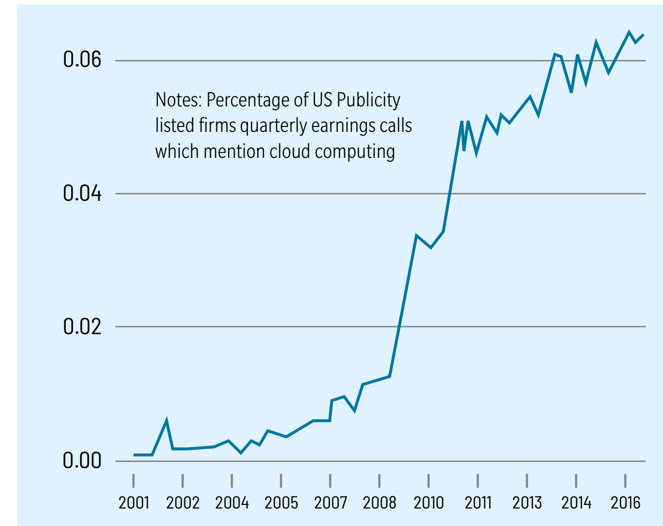

***Figure 5:** Cloud computing mentions in company earnings calls also exploded after 2006, and have also been flattening since 2016*

***Table 2:** Worldwide public cloud service revenue forecast (billions of US dollars)*

	2017	2018	2019	2020	2021
Cloud Business Process Services (BPaaS)	42.2	46.6	50.3	54.1	58.1
Cloud Application Infrastructure Services (PaaS)	11.9	15.2	18.8	23.0	27.7
Cloud Application Services (SaaS)	58.8	72.2	85.1	98.9	113.1
Cloud Management and Security Services	8.7	10.7	12.5	14.4	16.3
Cloud System Infrastructure Services (IaaS)	23.6	31.0	39.5	49.9	63.0
Total Market	**145.3**	**175.8**	**206.2**	**240.3**	**278.3**

BPaaS = business process as a service; IaaS = infrastructure as a service; PaaS = platform as a service; SaaS = software as a service
Note: Totals may not add up due to rounding.

Source: Gartner (August 2018) [3]

FIND OUT MORE

Research by Nicholas on a range of topics such as innovation, management and IT, including his 2018 paper 'Innovation, Reallocation, and Growth', is available to view at:
https://nbloom.people.stanford.edu/research

(2) https://www.gartner.com/en/newsroom/press-releases/2018-09-18-gartner-says-28-percent-of-spending-in-key-IT-segments-will-shift-to-the-cloud-by-2022
(3) https://www.gartner.com/en/newsroom/press-releases/2018-09-12-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2019

Innovation is slowing down

by Nick Bloom

(Stanford)

***W**hat explains the decline in productivity growth across the developed world? How will this impact jobs already threatened by the advance of robotics and artificial intelligence? Professor of economics at Stanford since 2005, Nicholas Bloom argues that growth is being curtailed as rich countries struggle to match their previous success as prolific innovators. Despite soaring investment in R&D, he says, the 'ideas productivity' of individual researchers has been in freefall for decades.*

Productivity growth is slowing down in the US *(see Figure 1)* and other developed countries *(see Figure 2)*. In the 1950s American productivity was rising by more than 3% a year. This period of incredible progress was driven by the rapid expansion of research universities like Harvard, MIT, and Stanford; research labs in firms like General Electric and Ford; and the commercialization of technologies developed in World War II. By the 1980s, however, productivity growth had halved to 2%. It has now fallen to just 1% per year.

This slowdown has sparked a debate among economists over the sources of the problem. Are statisticians underestimating output? Is the US mired in "secular stagnation", a prolonged period of low economic growth caused by insufficient investment? Or are recent innovations simply not as productive as those of the past?

In research with three fellow economists (Chad Jones and Mike Webb, Stanford; and John Van Reenen, MIT), I argue that ideas productivity - the productivity of science and discovery - has been falling for decades. Scientific discoveries and technical advances are getting harder to find, and spending on R&D has not been increasing fast enough to offset these declines in productivity. Innovation is slowing down.

Running to stand still

The creation of ideas is central to economic growth. This is driven by two things: the number of researchers (scientists and engineers) and the productivity of these individuals (ideas per researcher). Our analysis found that while there are a rising number of researchers, each one is becoming less productive over time *(see Figure 3)*. R&D efforts have been rising steeply for decades, but research productivity - the number of ideas being produced per researcher -has fallen rapidly.

WHERE'S THE DIGITAL REVOLUTION?

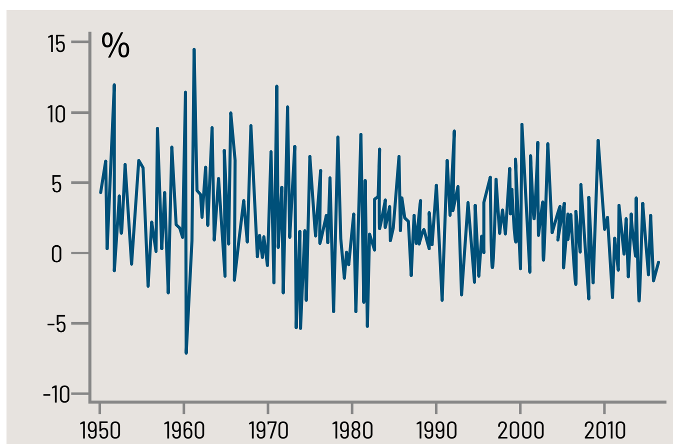

***Figure 1:** US productivity growth has been poor in recent decades*

GLOBAL SLOWDOWN

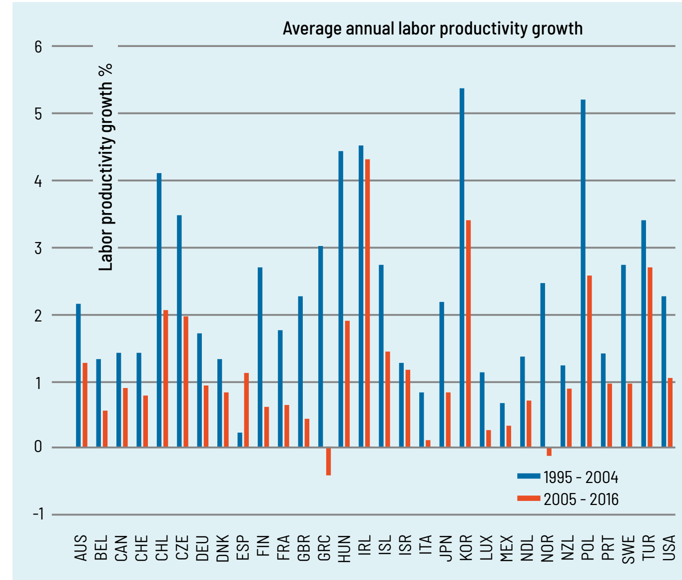

***Figure 2:** Productivity growth is slowing across almost all countries (including Singapore)*

Our analysis revealed that more than 20 times as many Americans are engaged in R&D today compared to 1930, yet their average productivity has dropped by a factor of more than 40. The only way the US has been able to maintain even its current lackluster GDP growth rate has been to throw more scientists and engineers at research problems. The US economy has had to double its research efforts every 13 years just to sustain the same overall rate of economic growth.

The acceleration of computer processing power is a telling example. Named after the co-founder of the computer chip giant Intel, Moore's Law holds that the transistor density of silicon chips will double roughly every two years.

RUNNING OUT OF IDEAS?

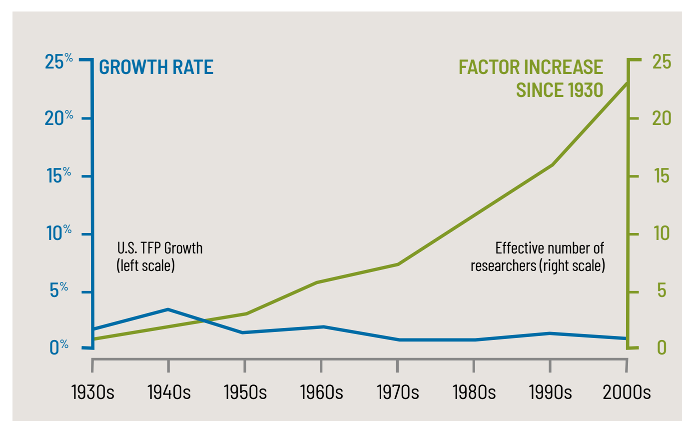

***Figure 3:** Productivity growth has been declining since 1940, while R&D spending is increasing. I argue that "ideas are getting harder to find"*

Such advances have enabled the creation of ever more powerful computers, which have transformed modern society. But maintaining that regular doubling today requires more than 18 times as many researchers than were needed in the early 1970s.

A similar pattern shows up in agricultural and pharmaceutical industries. For agricultural yields between 1970 and 2007, research effort went up by a factor of two, while research productivity declined by a factor of four, at an annual rate of 3.7%. In the pharmaceutical sector, research effort went up by a factor of nine between 1970 and 2014, while research productivity declined by a factor of five, at an annual rate of 3.5%.

> **"Are statisticians underestimating output? Is the US mired in "secular stagnation", a prolonged period of low economic growth caused by insufficient investment? Or are recent innovations simply not as productive as those of the past?**

We also examined the track records of more than 15,000 US public firms between 1980 and 2015. We found that even as spending on R&D rose, a vast majority of the firms experienced rapid declines in ideas productivity. The average firm now needs 15 times as many researchers as it did 30 years ago to produce the same rate of growth.

The apple tree model

So why has the productivity of scientists and engineers fallen so much? One explanation is that the low-hanging fruit of ideas have been plucked. To explain this 'apple tree model' of growth, we should travel back to the start of the Industrial Revolution in England.

Before 1750, productivity growth was close to zero. Most of the population in 1700 still worked on farms and were not much more productive than their ancestors under the Romans 2,000 years before. But from the late 1700s until about 1950, productivity growth began to accelerate. This is the era of "standing on the shoulders of giants". Each new invention - the steam engine, electric lighting, penicillin, and so on - made future inventors more productive. Growth took off as firms started creating industrial R&D labs, starting with those of Thomas Edison in 1876, while universities began to focus more on science and engineering research.

By 1950, however, the tide began to turn. The US reached its peak productivity growth of around 4% per year before the third phase of the 'apple tree model' began to set in. Humanity had made many of the quickest discoveries, and now unearthing new scientific truths started getting harder.

> The robot jobpocalypse is nowhere to be seen. One reason is that the types of low-skilled jobs that robots would replace have seen falling wages. So low-skilled workers have been forced to accept lower wages to price themselves into jobs

JOBS FOR HUMANS

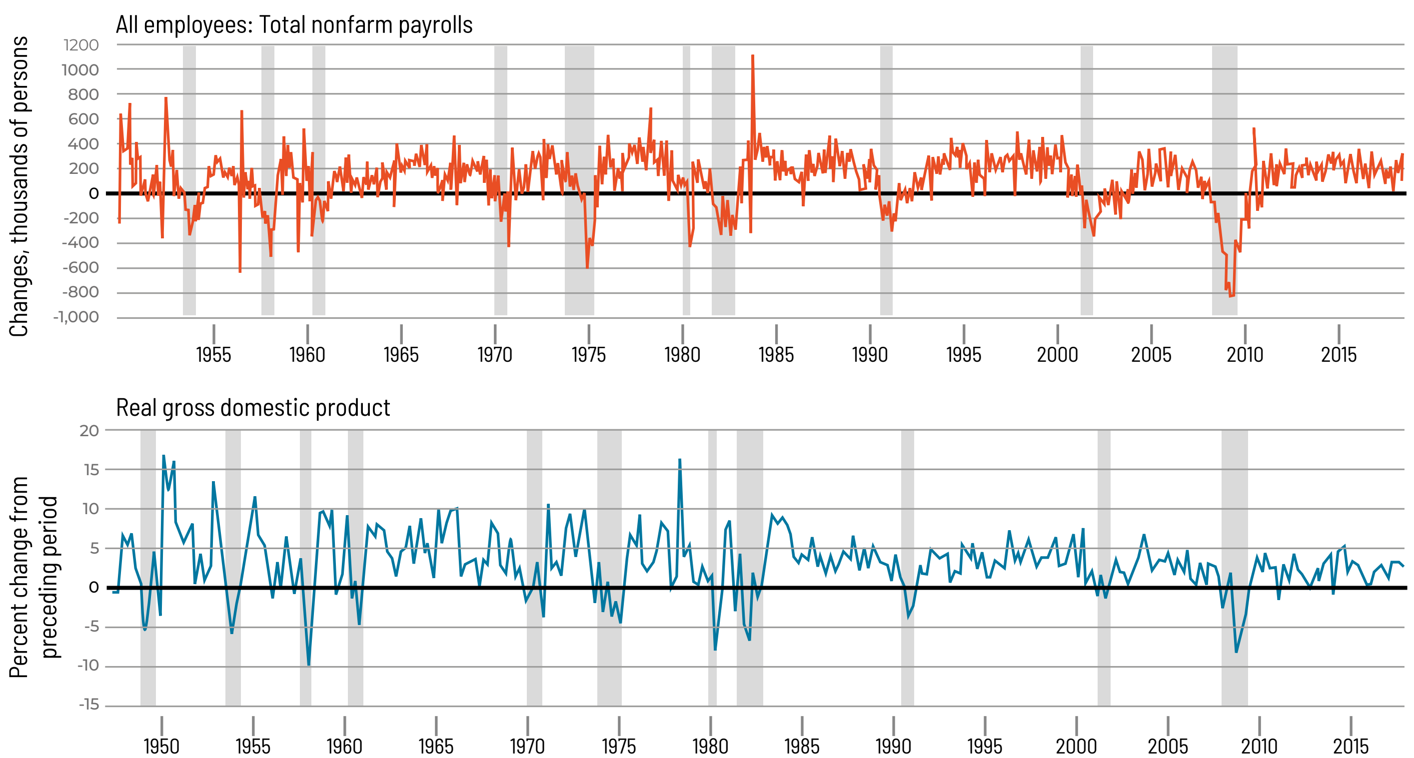

Figure 4: *US productivity growth is low because output growth is low and employment growth is high - this means robots are not (yet) stealing all our jobs*

What about the future?

In an accounting sense, productivity growth in the US has been slow because output growth has been low while employment growth has been high *(see Figure 4)*. In other words, robots have not been stealing our jobs (on net). In fact, the US has just experienced the longest and largest stretch of job creation since World War II (before which we lack accurate national accounts data). So the robot jobpocalypse is nowhere to be seen. One reason for this is that the types of low-skilled jobs that robots would replace have seen falling wages. So low-skilled workers have effectively been forced to accept lower wages to price themselves into jobs *(see Figure 5)*.

So I suggest two predictions:

A) The next 10 to 20 years are likely to herald similar trends of 1% productivity growth as the recent past. Compared to the 1950s to 1990s that is slow, but compared to the longer sweep of history going back 2,000 years it is still a blistering rate of advance.

B) In the longer-run artificial intelligence will be the great leveler - it has the potential to replace some types of higher-paid managerial and cognitive jobs that until recently have been safe from automation.

These higher-paid managerial jobs are also relatively much more expensive so the returns to developing technologies to replace them are greater. As such, artificial intelligence could reduce inequality and herald massive savings for companies employing large cohorts of graduate employees. But for these graduates - including the authors of this piece - it is less clear how beneficial that will be.

FIND OUT MORE

Research by Nicholas on a range of topics such as innovation, management and IT, including his 2017 working paper 'Are Ideas Getting Harder to Find?', is available to view at:

https://nbloom.people.stanford.edu/research

STAGNANT WAGES

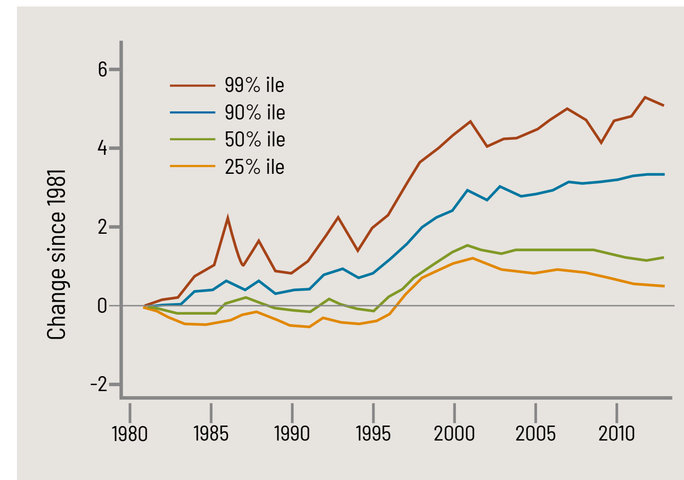

Figure 5: *Another barrier to the automation of the workplace is inequality: robots replace lower-skilled more manual jobs. But these manual jobs have seen stagnant wages*

TNIT News

TOULOUSE NETWORK FOR INFORMATION TECHNOLOGY

- *Scientific Director:* Jacques Crémer
- *Editorial contributions:* James Nash
- *Graphics:* Olivier Colombe
- *Illustrations:* I-Stock

TNIT
Manufacture des Tabacs
21 allée de Brienne
31015 Toulouse Cedex 6 France
+33(0)5 61 12 85 89

www.tse-fr.eu/digital

tnit@tse-fr.eu

The Toulouse Network for Information Technology (TNIT) is a research network funded by Microsoft, and coordinated by TSE. It aims at stimulating world-class research in the economics of information technology, intellectual property, software security, liability, and related topics.

All the opinions expressed in this newsletter are the personal opinions of the persons who express them, and do not necessarily reflect the opinions of Microsoft, TSE or any other institution.

Issue 20 June 2019

Toulouse School of Economics | Digital Center

Microsoft