TNIT News

TOULOUSE NETWORK FOR INFORMATION TECHNOLOGY

Innovation: an empirical approach

The art of management

Nicholas Bloom
(Stanford, NBER and TNIT)

Which discoveries should we patent?

Heidi Williams
(MIT, NBER and TNIT)

Issue 16 April 2017

Fostering innovation

Jacques CREMER & Priyanka TALIM

Innovation and technological change are at the heart of this issue of TNIT News, which features some of the latest research from TNIT members. Using data from more than 20,000 interviews with firms in 35 countries, Stanford economist **Nicholas Bloom** is one of the first researchers to conduct an empirical study of the relationships between information technology (IT), management practices and economic performance.

Bloom's research confirms that management practices are robustly linked to firm and national performance. Well-managed firms are more profitable, innovative and their staff have a better work-life balance. But it's not just businesses and their employees who stand to benefit: better management has also been linked to lower risk-adjusted heart attack mortality rates in hospitals and higher standardized test scores in schools.

What can policymakers do to surf this wave of potential? Bloom suggests improving management through increased competition and foreign investment, opening up trade and investing in workforce education. Encouragingly, the huge variation in management practices his work reveals suggests there is significant room for improvement around the world, such that adopting management best practices could lead to rapid employment and wage growth.

Also in this issue, **Heidi Williams** is an MIT micro-economist who examines how to foster groundbreaking scientific and medical research and protect intellectual property. Her work, which quantifies the impacts of "missing innovation" and earned her a MacArthur 'genius' award in 2015, has also been cited in US Supreme Court briefs.

For TNIT News, she reviews recent US Supreme Court rulings on patent eligibility and offers a progress report on her own quest to develop empirical evidence about the impact of gene patents on innovation. Williams' rigorous approach also has clear applications for research into patents for software and other technological advances.

For example, her study shows how datasets which are publicly available from the US Patent and Trademark Office can be combined with other data measuring innovation to develop quasi-experimental evidence that can inform business strategies, judicial decisions and public policies.

BYTES

TNIT member **Susan ATHEY,** an integral part of the network, published a very interesting article: *"Beyond Prediction: Using Big Data for Policy Problems"* in the February 3, 2017 issue of Science.

You can find more research by Susan on these topics at ***https://people.stanford.edu/athey/research*** and a survey article which explains some of the links between econometrics and machine learning in section 4 of her joint paper with Guido Imbens,

"The State of Applied Econometrics Causality and Policy Evaluation".

Congratulations to TNIT member **Daron ACEMOGLU** for winning the BBVA Foundation Frontiers of Knowledge Award for Economics, Finance, and Management, in honor of his prolific research contributions that have helped reshape his discipline.

Daron has published influential papers on the characteristics of industrial networks and their large-scale implications for economies. Working with a variety of collaborators, he has shown how economic shocks within particular industrial sectors can sometimes produce cascading effects that propagate through an entire economy.

The art of management

by Nicholas Bloom
(Stanford, NBER and TNIT)

More about N. Bloom

Well-managed firms have higher productivity, profits, market value and growth, says award-winning Stanford economist and TNIT member Nicholas Bloom. They're also more innovative, safer and cleaner. Information technology (IT) has often been claimed to be essential to effective management - encouraging billions of dollars to be spent on IT hardware and software - but it is only recently that researchers have been able to study the relationship between management, IT use and economic performance empirically.

To investigate such patterns, the World Management Survey (WMS) - a Harvard-LSE-Stanford research project - has collected firm-level management data across 35 countries since 2001. Using this survey tool, Bloom and his collaborators have shown management practices - defined in terms of monitoring, targets and incentives - are strongly related to development across countries and that effective management practices are strongly correlated with the intensive use of IT.

IT SOLUTIONS

Management is positively correlated with intensity of IT use

One of the major benefits claimed for the adoption of IT has been its ability to improve management systems. Comparing one basic measure of IT intensity against WMS management scores, Bloom found a clear strongly positive relationship, robust to a full range of controls for country, industry, firm size, age and other demographics. This positive relationship has been shown to be present across manufacturing establishments within the US. It is these types of modern management practices, says Bloom, that enabled US firms to drive the productivity miracle of the 1990s and 2000s.

MANAGEMENT MAGIC

Why some firms, industries and countries outperform their rivals

High-income countries with strong manufacturing sectors - such as Germany, Japan, Sweden and the US - tend to have the best management practices. Middle-income countries - such as Brazil, China and India - have moderate management practices, while those in Africa are really very poor.

Countries typically have management scores in line with their development - few countries are able to achieve substantial growth without improving their management practices. Increased productivity is usually necessary to afford the higher wages that come with development.

Within every country there is also a huge variation in management practices across firms. The wide dispersion of scores suggests that while it is possible to implement formal management practices across firms, they are not being implemented more widely. There are a number of well-managed firms - typically larger multinationals - operating in competitive markets with a highly educated workforce. But there are also poorly managed firms in most countries, typically family owned and managed, operating in protected and highly regulated sectors. Encouragingly, this spread across firms reveals there is significant room for improvement in management for many firms, that can lead to rapid employment and wage growth by adopting globally accepted management best practices.

When we consider three dimensions of management practices

> **Firms with higher management scores have substantially higher performance across every dimension analysed**

(monitoring, targets and people management), it is not surprising to see the top spot claimed by the US, the home of the business school and the plurality of the world's large multinationals, followed in second place by Japan, the birthplace of "lean manufacturing" (a modern management practice focused on data collection and continuous evaluation).

Within manufacturing, some industries tend to be managed better or worse. High-tech electronics and machine industries seem to be better managed. Lower-tech industries, such as furniture, leather and apparel, have lower management scores.

To what extent do better management practices translate into better economic performance? At the macro level, countries with higher management scores have higher GDP.

Meanwhile, firms with higher management scores have substantially higher performance across every dimension analysed: they are more productive, faster growing, more profitable and they have more employees, higher sales per employee and higher stock-market value. In other work looking at US firms, we also see that better managed firms are more innovative.

How robust are these results when controlling for factors such as industry, country and other potential performance-driving factors such as employee skills and firm size? We ran rigorous multivariate regression analysis of various indicators of performance on management scores, plus an increasingly complete set of controls.

Another concern may be that the management practices associated with economic success may differ in some countries. However, this survey evaluates very basic management practices which are generally perceived to be the core of good management practices. Furthermore, in a 2014 experiment, we re-estimated our performance regressions on the samples of firms in Asia, South America, Africa, Europe and North America, finding similar results in all regions.

One could also be concerned that this evidence is correlation based, so maybe rather than better management leading to superior performance, the relationship is reversed. Firms with superior economic performance might hire in good consulting firms to upgrade their management practices. In a 2013 study, we provided free management consulting to some large textile plants outside Mumbai, to help them adopt the modern practices measured by WMS, and compared their performance to another randomly chosen set of control plants. The adoption of these management practices took several months to occur, but eventually led to large increases in productivity.

Interestingly, the Indian experiment also found that firms were more likely to try to upgrade their management practices when facing tough times. If this type of endogeneity was common, it would lead to systematic underestimation of the impact of management on performance.

We also find that measures of work-life balance and family-friendly policies are positively correlated with WMS management measures, as are measures of worker safety such as the use of fire extinguishers and fire drills. By relentlessly focusing on cost minimization, better managed plants also reduce pollution and waste.

Management and IT

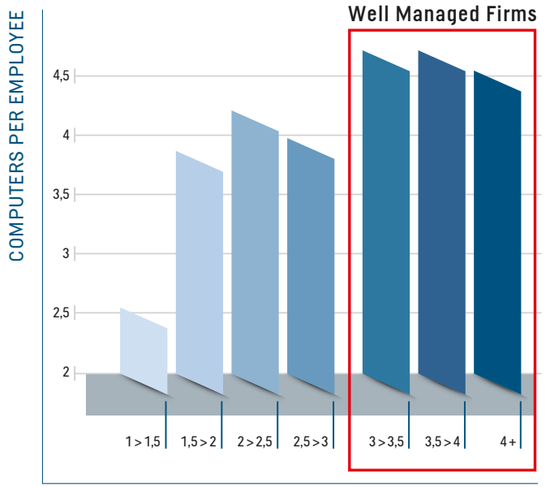

Note: Management scores from the WMS survey and computers per employee from the Harte Hanks establishment level IT survey. Data for 2884 manufacturing establishments across 10 countries from 2004-2009.

TNIT News
TOULOUSE NETWORK FOR INFORMATION TECHNOLOGY

WHAT WORKS

How to make better managers

Why is there so much variation in the quality of management practices across firms? Here are five drivers with direct policy implications.

1. Competition: Tougher product market competition is strongly related to better management practices. Competition forces badly managed firms to improve or exit the market. It also provides firms with lots of rivals to copy and learn from. Hence, a clear policy tool to improve management practices is increased competition - enabling firms to enter, removing any regulatory barriers on trade, FDI or market entry and vigorously policing anti-trust.

2. Ownership: Founder/family owned and managed firms tend to be managed significantly worse. The main issue here is not ownership, but control. Founder/family firms that have a founder/family member as CEO have low management scores, but founder/family firms with an external (non-family) CEO are just as good as other privately owned firms. Three policy levers can help here. First, minority shareholder protection to allow firm owners to hire professional managers without fear of expropriation. Second, an improvement in the general rule-of-law so that family owners can trust outside managers in their firms. Finally, FDI can play a critical role in spreading modern management practices. Foreign-owned firms are significantly better managed than domestic firms.

3. Regulation: Countries with lower regulation - as evaluated by the World Bank's Doing Business Index - have significantly stronger management practices. This link is likely to be driven by greater competition (international trade is highly competitive) and decision-making flexibility (allowing managers - rather than governments - to determine practices is likely to be more efficient).

> **Countries with lower regulation have significantly stronger management practices**

4. Skills/education: We found a strong relationship between the share of managers and workers with college degrees and quality of management. This makes sense when considering the importance of not just knowledge, but also implementation, of these best practices. Cultural changes in companies are only successful when there is significant buy-in from the employees, and this is often easier when workers are well educated and can be included in discussions about changes.

Management scores across countries

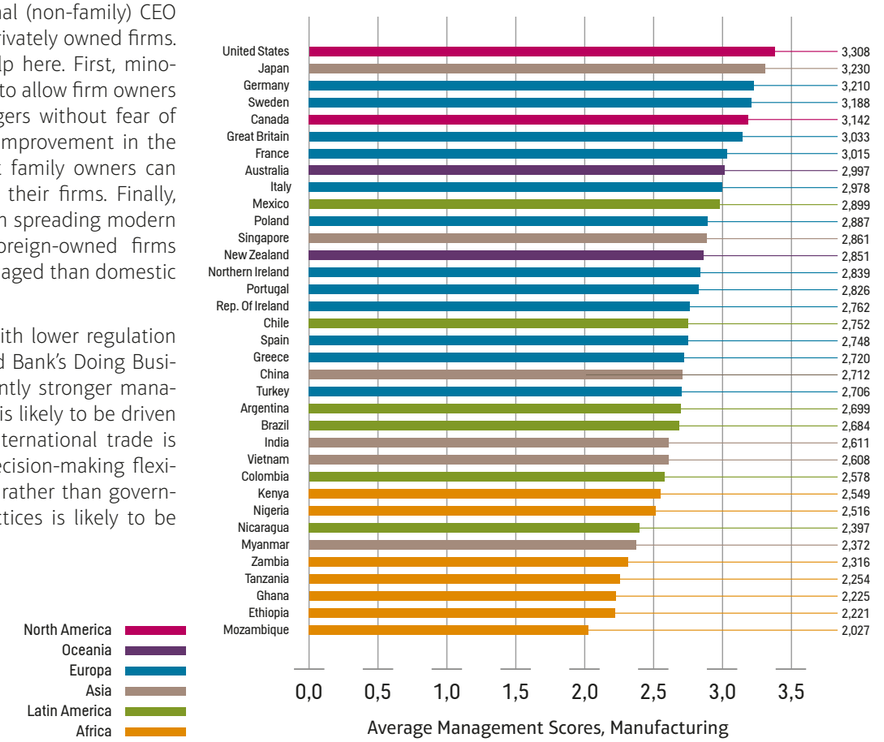

So creating incentives for continuing education of managers as well as employees is another policy action point. We can see human capital formation as a longer-term policy strategy in the general sense of the clear benefits accrued from a more educated populace. There are also short-term action points, such as identifying the skills most needed in each sector and offering training and workshops to current managers and employees. Training programs for basic operations across all sectors, such as inventory and quality control for manufacturing, could be a good place to start.

5. Information within and across countries: Across countries, the response to the question "How well managed is your firm?" is totally uncorrelated with the actual management score. Firms in countries with poor management practices do not seem to think their management practices are poor - if anything, they rank themselves more highly. But when we evaluate this self-perception across firms within the same country, we see a positive correlation, as better managed firms appear to rate themselves somewhat more highly compared to their local competitors. How can we explain these two apparently contradictory results?

Firms are able to evaluate themselves to a limited extent against their local competitors, but not against international competitors. Information on management practices seems to flow locally to some extent, but not at all across countries. A clear policy implication is that governments should try to expose domestic firms to the management practices of successful multinationals.

WHAT'S THE SCORE?

Measuring best practices

The WMS research team uses an interview-based evaluation tool that scores a set of 18 key management practices from one ("worst practice") to five ("best practice"). It has focused primarily on collecting data for more than 11,300 manufacturing firms, but it has also surveyed nearly 1,200 retail firms, 1,700 hospitals and 1,900 schools, finding similar patterns.

The WMS attempts to measure management practices in three key areas:

- *Monitoring:* How well do firms monitor what goes on inside the firm, and use this information for continuous improvement?
- *Targets:* Do firms set the right targets, track the right outcomes, and take appropriate action if the two are inconsistent?
- *Incentives/people management:* Are firms promoting and rewarding employees based on performance, prioritizing careful hiring, and trying to keep their best employees?

A firm earns a low score if it fails to track performance, has no effective targets, does not take ability and effort into account when deciding on promotions and has no system to address persistent employee underperformance. In contrast, a high-scoring organization frequently monitors and tries to improve its processes, sets comprehensive and stretching targets, promotes high-performing employees and retrains, rotates or exits underperforming employees.

To collect the data, graduate students who had some business experience were hired and trained from top US or European universities. These students interviewed plant managers who were senior enough to have an overview of management practices but not so senior as to be detached from day-to-day operations. Interviews were conducted in the managers' native languages.

Management and ownership

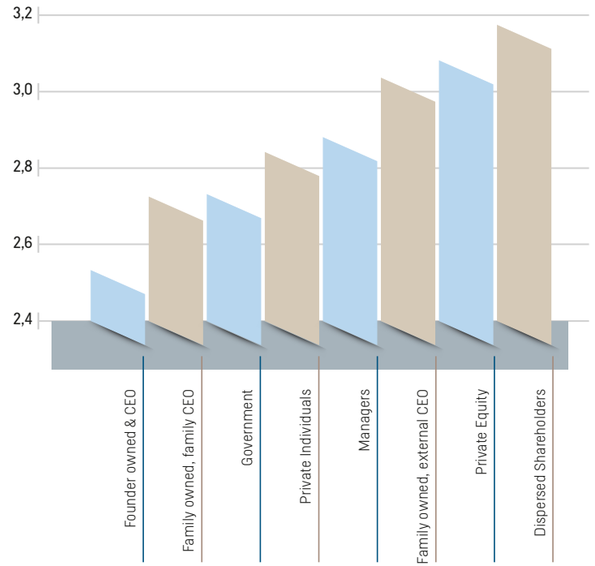

Note: Show the management scores for the 15,454 interviews across the countries in the WMS survey. Management is scored on a 1 to 5 basis for 18 questions. One ownership category called "Other" is omitted which contains 660 firms with a wide variety of ownership forms like employee coops, foundations and unknown.

These managers were not told they were being scored. They were told only that they were being "interviewed about their day-to-day management practices". To do this, we asked open-ended questions such as, "Could you please tell me about how you monitor your production process?" rather than "Do you monitor your production daily [yes/no]?". In this double-blind approach, interviewers were not told in advance anything about the firm's performance.

A variety of procedures were used to obtain a high success rate and to remove potential sources of bias. To obtain performance or financial data, only independent sources such as company accounts were used. A series of "noise controls" were also collected on the interview process itself (such as the time of day), characteristics of the interviewee (such as tenure in firm), and the identity of the interviewer.

For almost three quarters of all interviews, there was a second person listening in on a phone extension as a "silent monitor" to independently score the interview. For these double-scored interviews, the correlation across scores was 0.887, which shows that the two interviewers typically gave the same score.

Repeat interviews were also conducted at 222 firms from the manufacturing sample, using a different interviewer and a second plant manager within the same firm. The correlation between the first and second interview scores was 0.51 (p-value< 0.001). Part of this difference is likely to be real internal variations; no two plants will have identical management practices. The rest of this difference reflects measurement error.

Survey sample

The survey randomly sampled manufacturing firms with between 50 and 5,000 employees. The upper bound of 5,000 was introduced to exclude firms that are too large for us to proxy for the management of their firm from just one plant interview. Nonetheless, in other work on US manufacturers we find similar results extending the size distribution across the whole population.

The WMS tried to cover a range of major as well as smaller economies around the world across all six continents. Firms in the sample typically had around 250 employees, so cross-country comparisons are typically not impacted by large differences in firm size. Interestingly, developing-country firms tend to be relatively young, likely reflecting their rapid recent growth.

> **"Effective management practices are strongly correlated with the intensive use of information technology.**

KEY POINTS

- Management practices are strongly related to development across countries. Most firms in the US, northern Europe and Japan have world-class management practices. Their firms are: (I) continuously collecting and evaluating data to improve production efficiently, (II) setting stretching targets to motivate their employees to excel, (III) rewarding high-performing employees with bonuses and promotions (IV) and retraining or exiting underperforming employees. Firms in middle - and lower-income countries such as Brazil, China and India have substantially less advanced practices.
- Effective management practices are strongly correlated with the intensive use of information technology. One likely reason is that collecting and processing performance data - a key part of modern management practices - is facilitated by efficient IT systems.
- The key policy tools to improve management practices - and thereby raise employment, wages and growth - are competition, FDI, trade openness and workforce education.

Which discoveries should we patent?

by Heidi Williams
(MIT, NBER and TNIT)

More about H. Williams

Patents may be good or bad for innovation, but the hard evidence needed for effective policy and judicial decisions is often hard to come by. Winner of a MacArthur 'genius' grant, MIT's Heidi Williams is a specialist in the economics of gene sequencing, patent design and innovation incentives. Here, she reviews recent US Supreme Court rulings on patent eligibility and presents her rigorous empirical research on the question: do patents impede or encourage innovation? Williams' current focus is on genetic discoveries, but her work offers a template for future research on the value of patents in other fast-moving areas such as software development and financial services.

The informal narrative generally used to describe the requirements for obtaining a patent is that an inventor must submit a patent application that discloses an invention which is novel, non-obvious, and useful. In practice, there is an additional requirement: the invention must be "patent-eligible".

Section 101 of Title 35 of the US Code defines subject matter eligibility for patentability as follows: *"Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor…"*. In practice, the US Supreme Court has long interpreted patent eligibility as excluding abstract ideas, natural phenomena, and laws of nature. While it was not the first case to do so, as one example: the US Supreme Court's opinion on Diamond v. Chakrabarty clearly set out these boundaries. In that case, a genetic engineer (Chakrabarty) working for General Electric had developed a bacteria capable of breaking down crude oil, which he proposed as useful for cleaning up oil spills. The Supreme Court ruled in Chakrabarty's favor, arguing *"While laws of nature, physical phenomena, and abstract ideas are not patentable, respondent's claim is not to a hitherto unknown natural phenomenon, but to a nonnaturally occurring manufacture or composition of matter…"*

US SUPREME COURT RULINGS

A delicate balancing act

Recently, a set of four US Supreme Court rulings has clarified and - arguably - expanded the boundaries of what will be interpreted as non-patent eligible going forward:

1) **In Bilski v. Kappos** the Court invalidated patent claims on an investment strategy, announcing it supported a "high enough bar" on patenting abstract ideas that it would not *"put a chill on creative endeavor and dynamic change."* The claimed invention in this case was a system for buyers and sellers in energy markets to hedge against the risk of price changes; the patent application included a claim over a mathematical formula that can be applied to minimize risks from market fluctuations. The Court ruled that this claim was an abstract idea and therefore patent-ineligible.

> **For certain technologies, the social costs of patents may outweigh the social benefits**

2) **In Mayo v. Prometheus,** the Court invalidated patent claims on methods of using genetic variation to guide pharmaceutical dosing, expressing concern that *"patent law not inhibit further discovery by improperly tying up the future use of laws of nature."* The claimed invention in this case was a process by which physicians could measure patients' metabolite levels in order to determine the risk of thiopurine drug administration. The Court held that the relationship between metabolite concentrations and thiopurine risk is a law of nature and therefore patent-ineligible. Moreover, the Court argued that the steps outlined in the patent claims for administering and reading the test did not add anything that transformed the process into patentable subject matter.

3) **In AMP v. Myriad** the Court ruled to invalidate a subset of Myriad's gene patent claims, arguing that such patents *"would `tie up'…[genes] and…inhibit future innovation premised upon them."* The claimed invention on this case was human genes correlated with risks of breast and ovarian cancer. One technical detail that is critical to understanding the AMP v. Myriad case is that two types of nucleotide sequences were at issue: naturally occurring genomic DNA (gDNA), and complementary DNA (cDNA), the latter of which is produced in a laboratory using gDNA as a template. The Court decision drew a distinction between these two types of sequences: *"A naturally occurring DNA segment is a product of nature and not patent eligible…but cDNA is patent eligible because it is not naturally occurring."*

4) **In Alice Corp v. CLS Bank** the Court invalidated patent claims based on similar arguments. The claimed invention in this case was a scheme to mitigate "settlement risk," or the risk that one side to a settlement agreement will not meet their obligations.

The patent holder claimed a method for exchanging financial obligations, as well as a generic computer system and code to carry out the obligations. The Court held that this method of exchanging financial obligations was an abstract idea and therefore not patent eligible. They further found that the additional claims, which tie the method to the use of generic computer systems, did nothing to transform the idea into patent eligible subject matter.

LEGAL ANALYSIS

What is eligible for patent?

Taken at face value, the implications of these four rulings are incredibly broad. The Court has taken four technological areas - business methods, medical diagnostics, DNA, and software - and deemed that broad sets of inventions in those areas should no longer be eligible for patent protection. With the caveat that my background is in economics, not in law, my interpretation of these rulings is that the Court is relying on the Section 101 patent eligibility definition as a way of "carving out" certain technologies where they feel that the social costs of patents outweigh the social benefits.

For example, the Court's Mayo v. Prometheus decision argued that the patenting of abstract ideas may tend to impede innovation more than it encourages it. This is of course a theoretical possibility, and if such decisions were citing or otherwise building on rigorous empirical evidence that the social benefits of patents in those areas were indeed outweighed by the social costs, I could see the logical case for such decisions. However,

my read of the available empirical evidence is that there is essentially no strong empirical foundation either for or against that view.

Legal scholars have analyzed how the Court's reasoning over patent eligibility under Section 101 has progressed over the course of these cases. For example, in Bilski v. Kappos the Federal Circuit ruled that subject matter was not patent eligible unless it passed the "machine-or-transformation" test, which requires that a process is not patentable unless it is tied to a particular machine or transforms an article to another state or thing. Lemley et al. (2011) as well as other scholars have argued that this test is flawed, and consistent with that view the Court later argued that while the machine-or-transformation test is a clue to patentability, it is not dispositive to the question. As a second example, in Alice Corp v. CLS Bank the Court used a two-step test derived from its reasoning in the Mayo v. Prometheus decision: the Court first determined whether the claims were directed at one of the excluded classes; if they are, then the Court next determined whether the claims contain an *"inventive concept"* that *"produces something more than an attempt to claim the prohibited subject matter"* (see Burk 2014).

While these legal analyses have been extremely valuable in examining the reasoning used by the Court, a valuable complement to such analysis is the development of empirical evidence on what I view as the key question underlying these rulings: do patents tend to impede innovation more than they encourage it for these technologies?

Below, I summarize some ongoing work that is starting to develop evidence on this question in one context - the AMP v. Myriad case of gene patents - to give a sense of the direction that I hope future empirical research will develop to investigate the economic questions in this area.

THE AMP V. MYRIAD CASE

Do gene patents impede follow-on innovation?

The private firm Myriad Genetics was granted patent rights on human genes correlated with risks of breast and ovarian cancer. In 2009, the American Civil Liberties Union (ACLU) and the Public Patent Foundation filed suit against Myriad, arguing that many of Myriad's patent claims were invalid on the basis that DNA should not be patentable. In June 2013 the US Supreme Court unanimously ruled to invalidate a subset of Myriad's gene patent claims, arguing that such patents *"would `tie up'…[genes] and…inhibit future innovation."* In terms of relevance to Section 101 in particular, the AMP v. Myriad decision argued that *"[g]roundbreaking, innovative, or even brilliant"* discoveries of natural phenomena should be patent-ineligible, because patents *"would `tie up' the use of such tools and thereby inhibit future innovation premised upon them."* As discussed by Rai and Cook-Deegan (2013), the Court decision essentially aimed to draw a line between patent-eligible and patent-ineligible discoveries based on the *"delicate balance"* between patents prospectively creating incentives for innovation and patent claims blocking follow-on innovation. In the end, as discussed above, the Court drew this line by ruling naturally occurring gDNA patent-ineligible, and non-naturally occurring cDNA patent-eligible. Numerous legal scholars have argued that the distinction between DNA and cDNA is *"puzzling and contradictory"* (Burk, 2013) given that *"both isolated sequences and cDNA…have identical informational content for purposes of protein coding"* (Golden et al., 2013); in interviews, patent attorneys have expressed similar confusion.

While - consistent with the Court's view - there has been widespread concern that patents on human genes may hinder follow-on innovation, as there was essentially no empirical evidence available to either support or refute that assertion. In a recent working paper, economist Bhaven Sampat and I set out to try to develop empirical evidence on whether patents on human genes have hindered follow-on innovation in practice.

WEIGHING THE EVIDENCE

How to protect innovation in genetics

Investigating how patents on human genes - or more generally, on other existing technologies - affect follow-on innovation requires addressing two key challenges. First, in most markets it is extremely difficult to measure follow-on innovation empirically. We have a sense that in many or most markets, innovation is cumulative in the sense that any given technology is often an input into subsequent technological change, enabling follow-on discoveries.

But measuring this enablement is very difficult in practice. Second, ideally we would have something akin to a randomized experiment, where some human genes were randomly patented and others were not, and then we could confidently attribute any difference in follow-on innovation across patented and non-patented genes to a causal effect of the gene patents on follow-on innovation. In practice, inventors choose which genes to file patent applications for, and the patent office chooses which patent applications to grant patents to. Both types of selection raise the concern that any measured differences in follow-on innovation across patented and non-patented genes could reflect the selection of which genes were patented, rather than the causal effect of patents on follow-on innovation.

Our study aims to circumvent both of these challenges. To address the first - measurement - challenge, we take advantage of the fact that US patent applications claiming intellectual property rights over human genes are required to disclose the exact DNA sequences being claimed in the text of the patent.

By applying standard bioinformatics methods, these DNA sequences can be linked to gene identifiers, and these gene identifiers can in turn be linked to a variety of medical and scientific databases measuring follow-on scientific research and product development related to the human genome.

For example, gene identifiers are linkable to scientific publications in the PubMed *(https://www.ncbi.nlm.nih.gov/pubmed/)* database cataloging publications in the biomedical literature, to some datasets cataloging clinical trials in progress by both public and private drug development research organizations, and to datasets cataloging the availability of gene-based medical diagnostic tests. From a measurement perspective, this linkage from patent applications to the "bench" (scientific research) and to the "patient" (in terms of commercialized or under-development medical technologies) is remarkably complete.

Because we observe these measures of follow-on scientific research and product development for all human genes, this data construction is sufficient to tabulate a preliminary answer to our question: do patented human genes have more or less follow-on innovation? It turns out patented genes have higher levels of follow-on innovation. However, indicative of the selection bias challenge described above, it turns out that genes that will be patented in the future have higher levels of follow-on innovation even in the years prior to when these genes are patented.

This suggests that selection bias is a major concern, and that in particular the direction of selection is that patented genes look like they had higher potential for follow-on innovation even in the absence of their patents. This highlights the need for a research methodology that addresses this type of selection.

To address this second - selection - challenge, we develop two new quasi-experimental methods for estimating the causal effect of gene patents on follow-on innovation. First, we present a simple comparison of follow-on innovation across genes claimed in accepted and rejected patent applications. This method is valid if, conditional on being included in a patent application, whether or not a gene is granted a patent is as good as random. Consistent with this assumption, we document that patented genes look similar - in years prior to the patents being granted - to genes that are included in patent applications but not granted patents. Second, we develop a novel approach for estimating a causal effect of patents on follow-on innovation that takes advantage of the "leniency" of the assigned patent examiner.

> **"Do patented human genes have more or less follow-on innovation?**

While patent examiners have a uniform mandate, prior research has documented that in practice this mandate appears to leave patent examiners with a fair amount of discretion. We leverage this across-examiner leniency variation together with the fact that patent applications are quasi-randomly assigned to examiners, conditional on some covariates such as the application year and technology type, to approximate the following thought experiment: two equally high-quality gene patent applications come into the US Patent and Trademark Office at the same time, but one is more likely to be granted a patent because it is assigned to a more lenient examiner.

In contrast with the basic tabulations described above, both of our quasi-experimental approaches suggest that gene patents have not had quantitatively important effects on either follow-on scientific research or on follow-on commercial investments. These conclusions speak against the Supreme Court's argument in AMP v. Myriad, and more generally suggest that - as best we can measure - gene patents may not have had negative effects on follow-on innovation.

LOOKING AHEAD

A guide for future empirical research

Of course, this work is still preliminary. I highlight our study largely as an example of the type of empirical research that I would like to see more of in the future. Software patents may be good or bad for innovation, but rather than having policy or judicial decisions about software patents be based on theories or ideologies, these decisions would instead ideally be based on rigorous empirical evidence.

Our study provides an example of how standard datasets on patent applications and granted patents - which are publicly available from the US Patent and Trademark Office - can be combined with other data measuring innovation to develop quasi-experimental evidence on such questions.

TNIT News

TOULOUSE NETWORK FOR INFORMATION TECHNOLOGY

- *Scientific Committee:* Jacques Cremer
- *Production Manager:* Priyanka Talim
- *Editorial contributions:* James Nash
- *Graphics:* Olivier Colombe
- *Photos:* I-Stock

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

www.idei.fr

tnit@tse-fr.eu

The Toulouse Network for Information Technology (TNIT) is a research network funded by Microsoft and managed by the Institut d'Economie Industrielle (IDEI). 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, IDEI or any other institution.

More about TNIT

Issue 16 April 2017

Toulouse School of Economics

IDEI Institut d'Économie Industrielle

Microsoft