Innovation: Market failures and public policies*

Kevin A. Bryan[†] Heidi L. Williams[‡]

August 28, 2021

Abstract

Innovation is central to long-run economic growth. This chapter summarizes the state of the literature on the economics of innovation, highlighting open policy questions. We first articulate the key market failures in markets for innovation, and then discuss how both scientific norms and market-oriented policies help overcome those market failures. We close by discussing recent work on the diffusion of inventions as well as on the links between innovation and inequality.

*We are grateful to many colleagues for constructive comments and suggestions, including Pierre Azoulay, Yongmin Chen, Florian Ederer, Nancy Gallini, Patrick Gaule, Danny Goroff, Dan Gross, Ryan Hill, Erik Hovenkamp, Michi Igami, Adam Jaffe, Ben Jones, Alice Kuegler, Lisa Larrimore Ouellette, Enrico Moretti, Jacob Moscona, Kyle Myers, Kartik Shastry, Tim Simcoe, Carolyn Stein, and Scott Stern. We are also grateful to many of our students for helping to improve this chapter, including Ryan Broll, Dante Domenella, Maya Durvasula, Helen Kissel, Tamri Matiashvili, Gideon Moore, Helena Roy, Maya Roy, and Ralph Skinner. Finally, we are very grateful to Adam Jaffe and Mark Steinmeyer for their encouragement to write this chapter.

[†]University of Toronto, Rotman School of Management

[‡]Stanford University and NBER

Contents

1 Introduction

Innovation is at the core of many fundamental economic problems. The sustained invention and diffusion of new technologies during the Industrial Revolution brought us from the Hobbesian life of man - "solitary, poor, nasty, brutish, and short" - to the prosperity of the modern era. Innovation that replaces dirty technologies will be a key tool for mitigating climate change. The spread of "Green Revolution" crops allowed poor countries to grow without fear of famine. The world's seven largest private companies as of December 2020 all produce products that had not been invented fifty years earlier.[1] Indeed, the growth potential of the modern economy is itself an innovation problem: are we facing a future of stagnation (Gordon, 2017) or is this angst merely a replay of historical worries that technological growth had come to an end (Mokyr, Vickers, and Ziebarth, 2015)?

Innovation is *the invention, development, and diffusion of new goods, services or production processes*.[2] That is, innovation is the study of how society expands its production possibilities frontier. Innovation is an economic problem because it depends on the active choices of agents who respond to incentives. Nonetheless, innovation was not always seen as a primarily economic concern. Prior to 1960, only 11 articles in the *American Economic Review*, the *Quarterly Journal of Economics* and *Econometrica* combined had ever featured "invention" or "innovation" in their title.[3] Although new ideas and their spread are a constant feature of human history (Nisbet, 1980), their main driver was long thought to be psychological, sociological, or simply the result of serendipity. For instance, sociologists in the early twentieth century were concerned with the process by which new ideas - some of which are, as Emerson said, "in the air" - are adopted, spread, and modified (Gilfillan, 1935; Ogburn, 1922). This literature was largely focused on investigating which social structures were relatively more amenable or hostile to the adoption of new ideas.

The economic study of innovation was inspired by mid-century developments in industrial practice, government policy, and economic theory. Prior to World War II, large-scale industrial research laboratories became an important source of new invention (e.g., Hounshell and Smith (1988)). Successful directed wartime science led to Vannevar Bush's "Science, The Endless Frontier" (Bush, 1945), focusing on the importance of basic research and the role of government policy in its development. Finally, the dynamic economic theory of Schumpeter (1942) argued that the creation and diffusion of new goods was a more fundamental economic problem than the static Neoclassical welfare analysis which held the technological frontier constant.[4]

In response to these developments, a 1951 conference "Quantitative Description of Technological Change" was held at Princeton and supported by the Social Science Research Council. The meeting convened around 60 people including economists, economic historians, sociologists and historians – among them Gerard Debreu,

[1] On the economic history of technical change during the Industrial Revolution, see Mokyr (2010). Acemoglu, Akcigit, Hanley, and Kerr (2016) discusses the tradeoffs between tax-based and innovation-based climate change policies. Alston and Pardey (2014) summarizes the evidence on the contribution of innovation to international agricultural development, and Gollin, Hansen, and Wingender (2021) estimates that the Green Revolution contributed $83 trillion to world GDP since 1965. The seven largest private companies by market capitalization as of December 2020 are Microsoft, Apple, Amazon, Alphabet (Google), Alibaba, Facebook, and Tencent (Wechat).

[2] Many other definitions have been proposed, notably Schumpeter (1939): "we will simply define innovation as the setting up of a new production function," with even mergers included in this definition. This is surely too broad for our purposes.

[3] These are Epstein (1926) on whether invention is mainly driven by the profit-seeking motives of firms, Merton (1935) using patent data to study the rise and fall of innovation by industry, Bloom (1946) on adoption of old technology versus new invention in Hicks' theory of induced innovation, Terborgh (1950) with a short note on socialist versus capitalist innovation, Maclaurin (1950) on the history of the radio industry, Solo (1951) arguing that innovation is contra Schumpeter a normal business practice, Brozen (1951) with a qualitative description of factors affecting R&D and imitation, Maclaurin (1953) describing different propensities to do pure science, invent, or innovate over time and space, a short qualitative note by Duesenberry (1956) on how demand shifts with new innovations in an industry, Brown (1957) with theory and evidence of an effect similar to Arrovian replacement driving machine tool innovation, and Ruttan (1959) on the links between Usher and Schumpeter's definition of innovation. There are of course important papers on this topic in other journals (e.g., Plant (1934)) as well as articles on innovation and invention which do not use those words in their title (e.g., Nelson (1959b)). See Nelson (1959a) for a survey of the literature at that time.

[4] See Stephan (2015b) for a discussion of the influence of Bush's proposals. Linking all of these ideas was the work of Schumpeter's colleague Abbott Usher, an economic historian, particularly Usher (1954).

Simon Kuznets, Wassily Leontief, Rupert Maclaurin, Jacob Schmookler, and Abbott Usher.[5] Godin (2008) argues that publication of the conference proceedings – while originally envisioned – was abandoned because "the papers [were] in most cases of a very exploratory character." At the time, data on government R&D expenditures were close to nonexistent, very few papers had used information from patents to study corporate innovation, and the link between these efforts and economic growth was unclear.

The urgency of making progress on unlocking the "black box" of innovation was clarified with Solow's blockbuster 1957 paper (Solow, 1957). Solow shows that the share of long-run economic growth unexplained by changes in capital and labor inputs, referred to as total factor productivity (TFP), is as high as 85 percent. By construction, TFP is an unmeasured residual variation in output that cannot be explained based on observable inputs. However, subsequent work adjusting for labor quality and capital utilization suggested that much of this "Solow residual" reflects technological progress. Figure 1 presents a descriptive plot of TFP and gross domestic product (GDP) over time, providing one illustration of their co-movement. If much of economic growth is due to technological change, and if much of technological change is the result of incentive-driven choices, the implication is immediate: we had better understand what factors affect those choices.

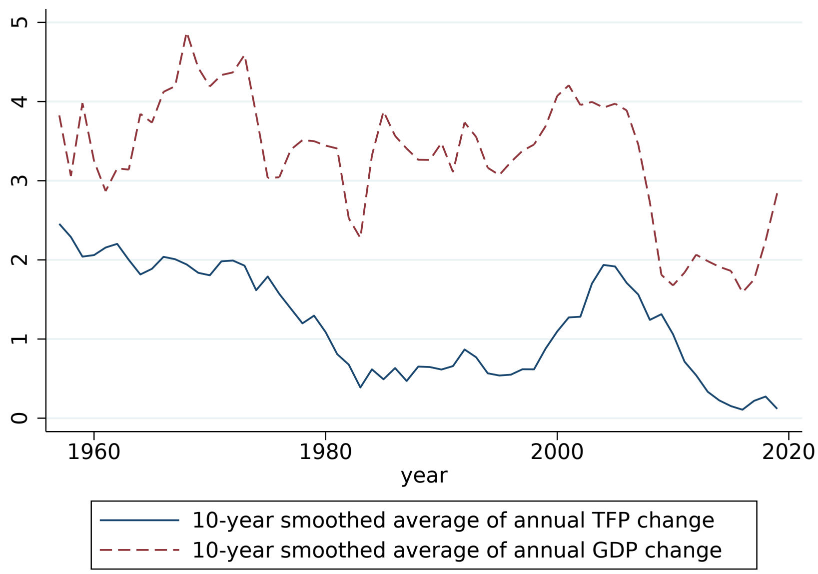

Figure 1: Changes in total factor productivity and GDP over time in the US

Notes: This figure displays 10-year smoothed averages of annual changes in total factor productivity (TFP) and gross domestic product (GDP). TFP here is the standard decomposition aside from adjusting for changes in labor and capital utilization (e.g., "labor hoarding" with shorter hours during recessions). Source: Fernald (2012).

Many questions immediately raise themselves. Why are certain inventions created when they are, where they are, by particular inventors? How do inventors choose the rate and direction of the research investments they pursue? How do inventors choose between improving existing products versus creating new ones? What role do public-sector scientists, tinkerers, and lone inventors play in innovation? What affects the speed at which these inventions flow to new consumers, new use cases, and new countries? What market structure leads to high levels of innovation? Do high labor or high capital costs shift effort? How does demand for an invention affect its production?[6] Is innovation optimally generated with laissez faire incentives, like normal goods? If not, why not? How does imperfect information about the value and difficulty of inventions matter? These core questions all appear in essays in the groundbreaking 1962 National Bureau of Economic Research

[5]The background in this paragraph is drawn from Godin (2008) and Godin (2012).

[6]On push versus pull theories of invention, see Scherer (1982).

(NBER) conference volume on the Rate and Direction of Inventive Activity (National Bureau of Economic Research, 1962) and related papers published by the economists who attended the event (e.g., Machlup (1962), Nelson (1959b), and Schmookler (1962, 1966)).

Despite the extensive research undertaken in the sixty years since the 1962 NBER Rate and Direction conference, many of these questions remain open. From a theoretical perspective, optimal policy often involves offsetting tradeoffs. Theory alone does not tell us whether innovation is under- or over-incentivized in laissez faire, or whether monopolies or competitive markets are more innovative, raising the value of empirical work. On the empirical side, measuring and valuing new ideas – the knowledge inventors build on, the goods which incorporate that knowledge, and the value of those goods to both firms and consumers – is very difficult. Moreover, isolating empirical variation which allows for the construction of compelling counterfactuals of how innovation might have been different with a different set of incentives is quite challenging.

The aim of this chapter is to provide a concise summary of the core problems in the economics of innovation. We begin with three fundamental concepts. First, what market failures exist in the production and diffusion of new goods - that is, why do we need an "economics of innovation" when we do not need an economics of cars or economics of apples? Second, how does inefficiency in the production of innovation harm economic growth? Finally, what makes empirically analyzing innovation policy questions challenging, and how have economists made progress in measurement and empirical analysis?

1.1 Market failures and innovation

The rate and direction of innovation responds to incentives (Schmookler, 1966; Schumpeter, 1942). This is of course true of all economic objects, the production of most of which we simply leave to the market. However, in an hugely influential article, Arrow (1962a) clarifies three fundamental market failures preventing laissez faire efficiency in the production of new ideas.

Arrow begins with the general equilibrium welfare theorems. Market equilibria are Pareto efficient as long as a set of well-known assumptions hold. Three of these assumptions are particularly important: the production transformation set should be convex,[7] the value of the good should not involve externalities, and production should be deterministic.

Production of the good called new knowledge - that is, innovation - violates all three of those assumptions. The first unit of a piece of knowledge is costly to create and subsequent pieces of that knowledge are free to replicate; ideas, as Thomas Jefferson noted, are like a flame being passed from candle to candle costlessly once the first candle is lit.[8] Whether one person or a thousand use a newly created idea, the cost of creating the idea remains the same, hence the production transformation set is not convex. This nonconvexity means market prices will not incentivize the right amount of innovation. Intuitively, efficiency requires that price equals the cost of the marginal unit of knowledge consumed, which is zero. But if the price of knowledge is zero, who will pay the fixed cost to create it first? If the inventor is the only one allowed to sell the good, how do we avoid the inefficiency of monopolies? That is, *the fact that initial inventors incur unique fixed costs means that innovation is either underappropriated or else the product of that innovation will be artificially scarce.*

Worse still for laissez faire efficiency, the social value of new ideas includes significant externalities. Ideas today are inputs into inventions tomorrow. There is no patent strong enough to cover ideas only *inspired* by the original invention, but the social value of the original invention surely ought to include that inspiration. Employee mobility is one source of this spillover, since firms are often unable to prevent their trained workers from leaving to new jobs. That is, *invention creates spillovers whose value is not incorporated into market prices.*

The fact that research inputs only produce knowledge with uncertainty is also problematic. In general,

[7] I.e., if I can produce 2 of good x using α of good y, then I need to be able to produce at least 1 of good x using $.5\alpha$ of good y.

[8] For an alternative historiography of the Jefferson analogy, arguing that his point was a natural rights and not an efficiency argument, see Sheff (2020).

the theory of Arrow-Debreu securities says that uncertainty does not harm economic efficiency as long as risks can be suitably hedged. With a market allowing proper hedging of risk, "the use of inputs, including human talents, in their most productive mode is not inhibited by unwillingness or inability to bear risks" (Arrow, 1962a). But Arrow, very much ahead of his time, noted that the uncertainty of research is partially related to the fact that we cannot observe the effort of the researcher. If a researcher gets a constant payout no matter whether she succeeds with her invention today or not, then surely she will shirk. The problem is then how to balance the efficiency-improving desire to shift risk away from individuals with the incentive problem of encouraging invention in the first place. *Invention involves uncertainties which cannot be fully hedged when they depend on unobservable effort by inventors.*[9]

Fixed costs borne only by the initial inventor, spillovers, and unhedged uncertainty collectively suggest that competitive markets will generate too little innovation. On the other hand, economic theory also suggests that research investments are in some contexts too high, for two reasons. First, inventors do not account for the fact that their R&D lowers the expected return of others trying to invent the same product (Loury, 1979; Reinganum, 1985). "Races" to enter the market first therefore occur. Second, a new invention may simply steal market share from other firms without generating any social benefit. For example, a new treatment for heart disease that is only a small improvement relative to existing drugs being sold might capture nearly the entire market even though the marginal social value of the new treatment could be relatively small. This is similar to the usual intuition that product entry can be excessive (e.g., Mankiw and Whinston (1986)).

But in practice, empirical analyses have suggested that – even after accounting for racing and business stealing effects – the social returns to R&D are much higher than the private returns, thus providing a justification for government-supported innovation policy.[10] Jones and Summers (forthcoming) notes that if all growth is due to domestic R&D, a baseline average social return is GDP growth divided by the social discount rate, over the share of GDP spent on R&D. The intuition is that a one-time investment in R&D today generates an infinite stream of higher production in the future. Using standard discount rate assumptions and historic US growth and R&D share figures, this baseline calculation suggests a social rate of return on R&D of up to 67%. After attempting to adjust for capital-embodied growth, health benefits unmeasured in GDP, and diffusion lags, among other factors, the average and marginal social returns of R&D nonetheless appear to be at least 20%. Bloom, Schankerman, and Van Reenen (2013) takes a different tack, using technological and market distance between firms to estimate the gross social returns to R&D. They likewise find large social returns, which are at least twice as high as the private returns.

As we will discuss, estimating the magnitude of this inefficiency – the gap between private and social returns – is incredibly challenging. Just as the magnitude of this inefficiency is difficult to estimate, many of the other most important parameters in innovation policy are also very hard to credibly uncover. We will discuss a number of modern empirical approaches that have made progress in pinning down these parameters in Sections 2 to 5.

1.2 Innovation and growth

There are good theoretical reasons, as Arrow has shown, to believe innovations are underproduced by the market. Worse yet are the consequences of this underproduction. Unlike many "standard" goods, innovations play a cumulative and fundamental dynamic role in economic growth. When innovation policy is suboptimal, growth is as well. This is true in exogenous growth models like Solow-Swan, in endogenous growth models, and in Neoschumpeterian models.

[9] The role of uncertainty comes up in many ways in innovation policy: uncertainty complicates contracting for the production of ideas when effort cannot be observed (Subsection 2.3), can cause holdup when there is uncertainty about opportunities for sequential inventions (Subsection 3.2.2), and can complicate targeting tax subsidies when distinguishing which corporate spending is actual R&D (Subsection 3.1). For a modern treatment of the importance of moral hazard, adverse selection, and private learning in generating inefficient innovation, see Halac, Kartik, and Liu (2016).

[10] See also Lucking, Bloom, and Van Reenen (2018) and Jones and Williams (1998).

Consider first the famous Solow-Swan model (Solow, 1956; Swan, 1956). Let the stock of labor grow exogenously, the savings and depreciation rates be a fixed proportion of output each period, and output be a Cobb-Douglas combination of capital, labor, and "technology" which scales up output for any level of capital and labor. In the short run society can become richer via capital deepening. However, once steady state capital per effective unit of labor has been reached, only technological growth can push the economy forward. Empirically, in line with the model, growth in capital and labor alone do a poor job of explaining the growing wealth of nations.

The Solow-Swan model assumes an exogenous rate of technological progress, in the sense of being independent of economic forces. Essentially, the model is driven by a technological "black box" with no microfoundations for what factors drive technological change. The natural next step, then, is to microfound this growth as the aggregated "endogenous" outcome of firm choice. Endogenous growth models where technology depends on firm R&D effectively build on two core ideas, one dating to Alfred Marshall, one to Edward Chamberlin.

The Marshallian idea is that ideas are nonrivalrous, so constant returns to scale technology for the firm production function of new inventions has increasing returns to scale at the aggregate level (Marshall, 1890). That is, one researcher may need a year to come up with a new product, but that new product can be used and recombined by other researchers without bound. The more ideas there are in the world, the more kindling there is for other inventors to spark (Weitzman, 1998).[11] The Chamberlin idea is that there exist markets without perfect competition. The quasirents that can be earned by successful innovators in imperfectly competitive markets can provide the incentive to do R&D (Chamberlin, 1962). Combining these two general ideas into a growth model that is both analytically tractable and which has nontrivial implications for innovation policy proved challenging.

The first successful attempt at incorporating increasing returns into a general equilibrium model comes from Romer (1986). Romer proposes an externality when firms use capital. Each firm's capital use increases overall knowledge in society via a learning-by-doing effect not captured by the profits of individual firms. This model, in the spirit of Marshall, generates a growth path with constant increases in capital, production, and technological growth.[12] This model's focus on the fact that knowledge is non-rivalrous – in the sense of being able to be used simultaneously by many economic actors – clarifies a key link between non-rivalry and increasing returns to scale. However, the model does not explicitly model knowledge production.

R&D explicitly appears in Romer (1990). Assume society can produce research on new production inputs, production of those inputs, or production of final consumption goods. Inventors of a production input have an infinitely-lived patent allowing them to earn quasirents from selling that good, hence incentivizing R&D. This model is canonical because it tractably allows analysis of how purposeful market invention can generate growth that is neither infinite nor zero in the long run. It is less satisfying because the form of competition among intermediate good producers (that is, inventors) is trivial: they are all monopolists with no strategic interaction.

While Solow-Swan has exogenous technology as the main driver of economic growth, and Romer focuses on the increasing returns to scale of non-rival technologies, the Neoschumpeterian models (Aghion and Howitt, 1992; Grossman and Helpman, 1991) focus attention on the competition between existing and new technologies. Invention today makes older varieties of the same technology less valuable. Therefore, whether a firm invests depends on whether it is destroying its own varieties (Arrow's "replacement effect") or destroying rival varieties. Further, the length of time a firm has market power for their inventions depends on the level of inventive effort other firms exert trying to create an improvement. Note that in exogenous, endogenous, and

[11] The growth models described in this chapter do not feature "path dependence," where random chance favoring inventions today can lead to large differences in outcomes tomorrow by changing the value of consumer adoption of goods with network effects, the ease of inventing follow-ons, and so on. See Subsection 4.2 for further discussion.

[12] There is a problem, however: the Romer model has a "scale effect" where population growth generates more capital use generates higher growth rates, counterfactually implying that growth is ever-increasing in population, eventually reaching a growth rate of infinity (Jones, 1999).

Neoschumpeterian models, economic growth is a function of the rate of innovation. In the latter two classes of models, policy around appropriation, invention costs, intellectual property, and market structure affect that innovation rate, and in turn affect economic growth.

These endogenous and Neoschumpeterian growth models make clear that long-run growth depends critically on the incentives for innovators. If the benefits of new goods are underappropriated, then we both have fewer things to consume today and we have fewer ideas to recombine into even better goods tomorrow. On the other hand, if inventors are given broad property rights over their inventions and their derivatives, then licensing frictions prevent today's silicon chips from becoming tomorrow's smartphones. With growth as the core long-run goal of innovation policy, a careful balance is needed between strong incentives today and overly-controlled inputs for tomorrow's inventors. Studying that balance has involved, as we will see, ever-deeper attempts to open the "black box" of innovation. Solow-Swan modeled growth as the result of technological change, Romer modeled technological change as the result of learning, and the models we examine later on in this chapter explicitly model the production function and strategic interaction inherent in that learning.

The fundamental link between economic growth and microeconomic innovation is increasingly being used to answer specific policy questions. For example, Aghion, Jones, and Jones (2019) discusses what needs to be true for automation and artificial intelligence advancements to lead to explosive economic growth. When human tasks are complements to automatable tasks, Baumol effects cause a growing share of the economy to be constrained by slow growth in the "human" sectors even as automation drops the price of many tasks to zero. Cavenaile, Celik, and Tian (2019) and Olmstead-Rumsey (2021) apply an endogenous quality ladder approach to studying the decline in productivity in the 2000s. The latter suggests that declining patent quality by laggard firms makes it harder to catch up to industry leaders, decreasing the laggard's incentives to invest in R&D, and causing market concentration to rise as productivity falls.

1.3 Empirical challenges

If innovation is inefficiently under-provided in private markets, there is a potential role for government policy to improve efficiency. Since efficiency in this market is first-order for long-run growth, getting policy right is critical. However, it is not enough for the policymaker to know that an inefficiency exists. Rather, the magnitude of that inefficiency, and the responsiveness of various margins being targeted as potential solutions, must be estimated.[13] For instance, what is the relative efficacy of public support for research and development as implemented through research grants versus tax subsidies? Are patents effective in inducing research investments – making a case for longer or broader patents – or are current patent laws too strong in the sense of hindering subsequent innovation? Would policies aimed at increasing competition, such as strengthened antitrust policies, increase or decrease innovation?

Answering these types of questions empirically is challenging for two key reasons. First, measuring the relevant economic phenomena in markets for innovation can be quite difficult.[14] Second, isolating sources of variation which identify compelling counterfactuals is particularly challenging.

1.3.1 Measurement challenges

Consider the problem of measuring knowledge spillovers. Knowledge spillovers are frequently cited as the central market failure justifying government intervention in markets for innovation. Yet, somewhat strikingly,

[13]These are – at least in a rough sense – the same two questions Arrow (1962a) argues were the two key questions of innovation policy: first, how shall the amount of public subsidies to R&D be determined? And second, how shall efficiency of their use be encouraged?

[14]The empirical literature on growth accounting led by Solow (1957) and others measures innovation by the share of long-run economic growth that could not be explained by changes in capital and labor inputs. This "residual" approach was helpful in highlighting the likely importance of technological change in driving economic growth, but left an important gap in terms of how to directly measure the rate and direction of technological innovation.

evidence for the existence and magnitude of this market failure is quite thin – we think in large part because of the challenge associated with measuring these spillovers. As Krugman (1991) argues, "*knowledge flows...are invisible; they leave no paper trail by which they may be measured and tracked.*"

A key leader in developing direct measures to tackle this issue was Zvi Griliches. While the idea that the rate and direction of technological innovation were influenced by economic incentives was not new, there were essentially no quantitative estimates of those relationships at the time Griliches began his career. While Griliches made a variety of important contributions, perhaps one of his most important efforts was working systematically with a variety of collaborators over several decades to construct newly digitized data on patent statistics – not out of an interest in the patent system per se, but rather out of an interest in the potential role of patent data as an economic indicator of innovative activity.[15] As articulated in his 1990 *Journal of Economic Literature* paper (Griliches, 1990): *"In this desert of data, patent statistics loom up as a mirage of wonderful plentitude and objectivity. They are available; they are by definition related to inventiveness, and they are based on what appears to be an objective and only slowly changing standard. No wonder that the idea that something might be learned from such data tends to be rediscovered in each generation."*

Griliches (1990) details an intellectual history of the development of some of the key early patent-related datasets, including the National Bureau of Economic Research (NBER) patent data (an effort that grew into the work of Hall, Jaffe, and Trajtenberg (2002)), work by F.M. Scherer (such as Scherer (1965)), the Yale and Carnegie Mellon groups (Levin et al. (1987) and the later work of Cohen, Nelson, and Walsh (2000)), and the Science Policy Research Unit group at the University of Sussex. From a modern-day perspective, these and related patent data efforts have in many ways been quite successful. For example, the Hall, Jaffe, and Trajtenberg (2002) linkage between the Compustat data and granted US patents has garnered over 4,000 citations across a wide variety of fields including economics, finance, law, management, and strategy. The US Patent and Trademark Office (USPTO) has greatly expanded access to their administrative data in recent years,[16] and independent researchers have also been developing novel techniques for digitizing more historical USPTO records (see, e.g. Berkes (2016)).[17]

Yet Griliches himself was also one of the key critics of reliance on patent data. For example, at the 1962 NBER conference (Griliches, 1962) he noted: *"[I]nventions may be the wrong unit of measurement. What we are really interested in is the stock of useful knowledge or information and the factors that determine its rate of growth. Inventions may represent only one aspect of that process and may be a misleading quantum at that."* More specifically, measuring innovation using patent statistics has a number of important limitations.[18] As emphasized in the survey work undertaken by the Yale and Carnegie Mellon groups (Cohen, Nelson, and Walsh, 2000; Levin et al., 1987), many inventions are not patented and – perhaps more concerningly – the propensity to patent a given invention appears to vary tremendously across industries. Hence, while researchers often see patents as having the advantage of being a standardized measure of innovation that is relevant across industries, in fact the selection of which and what types of inventions are patented in different industries is likely first-order. A related issue – as documented by Pakes (1986), Schankerman and Pakes (1986), and others – is that patents vary tremendously in their quality or value. We will discuss measures of patent quality that have been developed in Section 3.2, but simple counts of patents are in many contexts not obviously a meaningful measure. Other issues can also arise with patents – for example, firms sometimes explicitly tie incentive pay to patenting, again potentially changing the appropriate interpretation of patent

[15] Griliches is also credited with leading efforts to create linkages across key administrative datasets, and to make such data available to researchers through the US Census Bureau's regional census data centers.

[16] See, for example, https://www.patentsview.org and https://bulkdata.uspto.gov.

[17] For a recent overview of this and other large-scale historical patent data efforts, see Andrews (2021). The Innovation Information Initiative is compiling a variety of patent-related data resources here: https://iii.pubpub.org.

[18] A separate issue is that patent data themselves are often misused; see Lerner and Seru (2017). See also Hall, Jaffe, and Trajtenberg (2002), which documents a number of best practices on how to normalize patent data by time and technological area measures.

counts.[19]

Patent statistics are hence a useful proxy for innovation but also have important limitations, naturally raising the question of what alternative measures are available. Scientific publications are an alternative source of publicly disclosed data documenting scientific advances. However, the decision of whether to publish a given scientific finding critically depends on the norms and institutions in which that discovery is made. As we discuss in Section 2, academics face strong incentives to publish in order to establish priority of discovery and collect credit for their work. Discoveries made by non-academics, in contrast, are published less often and more selectively.[20] One could alternatively consider R&D investments as a proxy for innovation. Although fruitful in some contexts such as the work of Sakakibara and Branstetter (2001), firm-year aggregate measures of R&D investments – even in cases where they are publicly reported – in most cases provide little insight into the types and value of inventions being pursued by firms.[21]

In recent decades, perhaps the most progress has been made in the context of measuring technological innovation in health care markets, where direct measures of research investments – such as clinical trial starts and new drug approvals – are available and can be directly linked to the patient populations relevant to those technologies.[22] Of course, direct measurement of research investments in clinical trials does not directly translate into statements about changes in welfare, which depend on how new innovations impact prices and health outcomes, but opportunities to construct such direct linkages to welfare-relevant outcomes are quite rare.[23] Nonetheless, health care markets are a useful illustration of how the innovation process itself (here, the process of clinical trials and drug approvals) can sometimes generate data that encodes information which in many other domains would be challenging or impossible to observe directly; the data constructed by Sampat and Williams (2019) and Hill and Stein (2020a,b) are two other illustrations of this idea.

This measurement problem is even more challenging for other innovation questions. Consider as an example the problem of how to measure spillovers from new ideas. As articulated in Section 1.1, if one firm creates something truly innovative, that invention may create spillovers in the sense that other firms or individuals might either copy or learn something from the original research at less than the full R&D cost and without compensating the original inventor.

Returning to the Krugman (1991) quote: "*knowledge flows...are invisible; they leave no paper trail by which they may be measured and tracked.*"[24] An inventor may know they used a given assay to help develop a drug, but what written document would include the computer chip used to process the data? What of the theoretical solid state physics that made the computer chip possible? While the thoughts in inventors' heads are of course unobserved by the econometrician, Jaffe, Trajtenberg, and Henderson (1993) argue knowledge flows do sometimes leave a paper trail in the form of citations in patents – that is, acknowledgements of the use of knowledge in subsequent patents as recorded in so-called front page patent citations. Patent citations

[19] Pekari (1993) (in Finnish; see Toivanen and Väänänen (2012) footnote 2 for a discussion in English) provides qualitative evidence on this issue from case studies and interviews with sixteen actively patenting companies in Finland. In eleven of the sixteen companies (and in all of the large companies that were interviewed), there were explicit rules for rewarding inventor employees. In large companies, the reward structure is typically composed of three payments: one at the time of notice of invention, one at the time of patent grant, and one later payment as the value of the invention is revealed over time. Our informal sense is that many countries require by law that firms tie pay to patenting; for example, Lobel (2019) argues, "[a]n anomaly of the American legal system is the complete absence of any requirement for businesses to compensate their employed inventors."

[20] The work of Fiona Murray, Scott Stern, and others has highlighted that the process of publishing commercially valuable discoveries is often closely coordinated with the timing of patent filings, e.g. Murray and Stern (2007).

[21] As discussed in Section 3.1, concerns also arise with the level of reported R&D potentially being manipulated by firms due to incentives provided by policies such as R&D tax credits.

[22] See, for example, Acemoglu and Linn (2004), Finkelstein (2004), Kyle and McGahan (2012), Blume-Kohout and Sood (2013), Budish, Roin, and Williams (2015), Dubois, de Mouzon, Scott-Morton, and Seabright (2015), Sampat and Williams (2019), Yin (2008), and Yin (2009) on clinical trials as well as Williams (2013), Galasso and Schankerman (2015), and Sampat and Williams (2019) on medical diagnostic tests and other medical instruments.

[23] Budish, Roin, and Williams (2015) links technological innovation in cancer drugs directly to changes in health outcomes, but is unable to do a full evaluation of welfare even within that narrow setting.

[24] Although not directly relevant to our discussion in this section, Krugman (1991) continues, "*...and there is nothing to prevent the theorist from assuming anything about them that she likes.*"

serve a legal function, in the sense that applicants have a legal duty to disclose prior art when applying for a patent. However, patent examiners can also add citations – which could have a different interpretation than spillovers – and some citations may be internalized in the sense of being accompanied by a licensing payment, which would almost always be unobserved but which would mean the word “spillover” is less appropriate.[25] On net, patent citations are thus helpful but are thought to almost certainly be an incomplete metric of spillovers.[26]

If we want to look beyond direct patent citations, the trick as Griliches describes it is to define a dimension along which spillovers are mediated. Economists have long analyzed input-output matrices – that is, records of how industries trade with each other – as one way of parametrizing similarities and differences across firms. Indeed, early studies such as Brown and Conrad (1967) use input-output tables to construct matrix weights which assume spillovers are embodied in purchased inputs. However, it is unclear the extent to which input-output tables are relevant to knowledge spillovers. In a similar vein, studies such as Bernstein and Nadiri (1989) use the structure of industry codes (specifically, SIC codes) to construct assumptions about how firms would benefit from R&D done by other firms in related industries. However, other than defining “own industry” spillovers, SIC codes do not lend themselves to a natural ordering of which industries are more versus less similar. Griliches (1992) articulates this problem by asking whether the industry code for ‘leather’ should be closer to that for food or that for textiles.

The literature has primarily focused more recently on two dimensions of spillovers: technological distance and geographic distance. The classic reference on technological distance is Jaffe (1986), which uses an early version of the NBER patent data to define the technological position of a firm based on the USPTO-assigned technological classes in which the firm had previously patented. Given that parametrization, he defines a firm’s “potential spillover pool” as a weighted sum of other firms’ R&D in that technology space. Bloom, Schankerman, and Van Reenen (2013) builds on Jaffe’s work and develops a more flexible Mahalanobis extension. The classic reference on geographic distance is Jaffe, Trajtenberg, and Henderson (1993), which asks whether patent citations are geographically localized relative to a counterfactual geographic distribution of citations. While the details of this approach were subject to debate (Henderson, Jaffe, and Trajtenberg, 2005; Thompson and Fox-Kean, 2005), the key idea behind this approach has been incredibly influential. One central issue that arises in both approaches is the distinction between testing for the existence of spillovers versus quantifying the magnitude of spillovers: by construction, defining a dimension along which to search for spillovers will essentially always result in an underestimate given that total spillovers likely accumulate on multiple dimensions. Bloom, Schankerman, and Van Reenen (2013) is perhaps the leading paper in this literature, which tackles this issue in part by parameterizing connections across firms in multiple dimensions.

1.3.2 Inference challenges

As we have seen, measuring the relevant economic phenomena in markets for innovation can be quite challenging. Of course, in many cases researchers are interested not just in *constructing* measures of — say — inventive activity, but rather in using such measures as an input in studies aimed at understanding how a

[25] On examiner-added citations, see Alcácer, Gittelman, and Sampat (2009). The authors motivate their analysis with this quote from Gregory Aharonian, editor of the electronic newsletter PATNEWS: “...the economists who write about patent citation analysis have little experience with patent searching, and don’t realize how worthless most patent citations are for measuring anything. For example, many of them assume that the citations that appear on the front of the patent were all used and discovered by the inventor. They then use that assumption to measure flows of information between companies and inventors.. What they don’t realize is that many citations are found either by the examiner or by professional searchers .. so that such citations do not measure anything about information flow or patent importance.” Empirically, Alcácer, Gittelman, and Sampat (2009) documents that examiners added approximately 40% of all citations in US patents in 2001–2003. Bryan, Ozcan, and Sampat (2020) argues in-text citations also appear to be less frequently added by examiners. On licensing payments, see Arque-Castells and Spulber (2019) which documents that “spillovers” across firms are correlated with licensing in at least one dataset. Also related is the work of Roach and Cohen (2013) which argues that patent citations reflect knowledge flows from public research, but appear to miss knowledge flows that are more private and contract based in nature, as well as those flowing from privately-funded basic research.

[26] See also Jaffe, Trajtenberg, and Fogarty (2000), which draws a similar conclusion based on the results of a survey of inventors.

given policy shaped inventive activity. While some aspects of this type of policy analysis are common to empirical analyses in other fields, there are also unique challenges that arise in constructing counterfactuals in innovation markets.

Consider, as an example, estimating the effects of changes in state-specific tax rates on the cross-state mobility of US inventors. This type of question has frequently been analyzed in the public finance literature. The outcome — inventor mobility — is observable. The choice of an inventor in California of whether to stay in California, move to a lower-tax state, or move to a higher-tax state, can naturally be parametrized: the set of potential states an inventor could move to are known, and the tax rates relevant in each jurisdiction are observed. Moreover, the relevant behavioral elasticities can naturally be inferred by scaling the mobility response to a given policy change by the difference in tax rates across states. For all of these reasons, it is perhaps not surprising that cross-state and cross-country variation in tax rates have been leveraged very successfully in service of estimating mobility responses to taxation.

In contrast, consider attempting to estimate how changes in the US patent term — say, a change from 17 years to 20 years — affects inventive activity. Inventors from anywhere in the world can choose to file for a patent in the US market. Any "direct" change in inventive activity in response to this policy change would in expectation generate knowledge spillovers that — although they may be disproportionately local — could affect individuals globally as well. In such a setting, designing an empirical methodology that contrasts a treated group with a control group is conceptually quite challenging. This difficulty is all the more important given that measures of research investment vary across technologies over time for many reasons, some of which — such as the cost or ease of different technological opportunities — are inherently unobservable to the econometrician.

Unlike in the tax rate example in which the set of US states naturally provides a choice set of where inventors might consider moving, it is rare to have a map of "potential inventions" which can be used to directly estimate which inventions could have been developed but were not. Nagaraj and Stern (2020) makes a version of this argument, relating the ex ante known structure of the human genome as analyzed in Williams (2013) to a space onto which potential and actual inventions can be mapped.[27] We will discuss analogous challenges that arise in the various empirical contexts discussed throughout this chapter.

1.4 Overview of this chapter

We have seen that innovation involves substantial market failures, is fundamental to economic growth, and yet is difficult to study empirically. In the remainder of the chapter, let us turn to specifics by discussing the state of the literature and open questions. In Section 2, we consider how scientific norms and institutions complement and interact with market incentives. In Section 3, we examine a number of market incentives which either shift the payoff to inventors directly through tax or subsidy policies, change the nature of post-invention competition through patents or antitrust intervention, or change the nature of invention inputs through immigration or agglomeration policy. In Section 4, we discuss why inventions may diffuse slowly, how this can be measured, and when policy can speed this diffusion. Finally, Section 5 discusses the bidirectional link between inequality and innovation.

It is worth noting explicitly that this chapter focuses on public policy levers which can address specific market failures. Recent methodological advances in economics, many of which were first developed outside of innovation, are permitting particularly fruitful research in line with that focus. Nevertheless, the field of innovation, or more broadly innovation and entrepreneurship, without question tackles problems of greater scope that we can fully cover here. Indeed, the innovation chapter in a previous volume of this Handbook (Cohen and Levin, 1989) largely examines the classic Schumpeterian hypotheses of how innovation varies with firm size and market concentration, and the relation of these factors to markets for technology. It is therefore

[27] Budish, Roin, and Williams (2015) provides another example, where the structure of cancer — which can affect different parts of the body and be diagnosed at different stages of disease — provides a natural map onto which observed research investments can be mapped, and from which "missing" research investments can be inferred.

worth explicitly noting some topics we will not give their full due: markets for technology and the frictions in them (Arora, Fosfuri, and Gambardella, 2001; Arora and Gambardella, 2010), firm-level heterogeneity in innovative output and profitability (Cohen, 2010), creative works covered by copyright (Biasi and Moser, forthcoming; Giorcelli and Moser, 2020; Li, MacGarvie, and Moser, 2018), entrepreneurship and innovation finance (Gompers and Lerner, 1999; Hall and Lerner, 2010), innovation prizes (Kremer and Glennerster, 2004), evolutionary models of innovation (Dosi and Nelson, 2010; Nelson and Winter, 1985), user innovation (von Hippel, 2010), within-firm considerations (Teece, 2010), general purpose technologies (Bresnahan, 2010), and innovation systems (Soete, Verspagen, and ter Weel, 2010).

2 Science as a non-market incentive

Our primary focus in the remainder of this chapter will be on how the market failures articulated in Section 1.1 can be addressed via market-based policies which change prices (via taxes or subsidies) so as to improve the alignment of market rewards with the first best (see Section 3). However, before turning to discuss these policies, we first discuss one particularly important set of non-market incentives: namely, scientific norms and institutions.

Scientific exploration occurs partly for reasons of pure curiosity and creativity; the economist Josiah Stamp argued almost one hundred years ago that "[t]he sense of curiosity and the idea of fame play a greater part than the economic reward" (Stamp, 1929). However, beyond pure curiosity, scientists are motivated by academic tenure, scientific credit, or accolades like major prizes. And their ability to produce breakthroughs is affected by a number of policy mechanisms, such as government funding of basic scientific research. We will not here provide a comprehensive review of the economics of science literature (for recent reviews, see Stephan (2010, 2015a)). Rather, we will focus on a few concrete examples of topics in the economics of science which have seen recent advances and which seem ripe as topics for future work.

Two important caveats are needed. Science policy affects innovation in a much deeper way than just the production of "basic research." In the so-called linear model, basic research - defined by Bush (1945) as that "performed without thought of practical ends [resulting]...in general knowledge and an understanding of nature and its laws" - leads to applied research which is then iterated into commercial products. This model is generally not now seen as accurate. Instead, science includes a Pasteur's Quadrant, in the evocative framing of Stokes (1997), where commercial opportunities inspire fundamental research just as much as the reverse. Louis Pasteur's research was, as Stokes notes, explicitly applied in its focus, but the methods he pursued led to the field of microbiology. At the time of the Industrial Revolution, Britain was not a world leader in pure science, but did have a culture where "technologically valuable knowledge had penetrated into the productive layers of society" (Mokyr, 1999). Science policy is therefore not simply the production of fundamental early-stage research by a university researcher.[28] Instead, science policy affects the full gamut of innovation from basic to applied via the incentives to create, publish, and promulgate new ideas and inventions in the absence of profit from commercially-sold products.

A second caveat is that fundamental science can also be produced by private firms. For much of the twentieth century, industrial R&D labs proved especially important: nine Nobel Prizes were awarded based on research at Bell Labs alone. Rosenberg (1990) argues that firms perform pure science both because it sometimes leads to useful products and also because it provides firms with a base from which to draw on academic science. By 1940, the average number of forward citations to a scientific publication published by an industrial researcher exceeded that of the average paper published by an academic researcher (Arora, Belenzon, Patacconi, and Suh, 2019).

[28] Furthermore, science often draws on ideas that come from more applied research. This is both because science requires tools developed elsewhere and because researcher ideas often come from users. On the former, see Furman and Teodoridis (2020) on how the release of Microsoft Kinect motion sensing technology affected researcher trajectories. On the latter, von Hippel (1988) describes this pattern of "user innovation" across a number of industries; users may be particularly important for helping radical innovation in contexts where user demand is uncertain (Chatterji and Fabrizio, 2014).

That said, industrial science in large firms has declined substantially since the era of the industrial research laboratory in the early- and mid-20th century (Costa and Lamoreaux, 2011). Arora, Belenzon, and Patacconi (2018) shows that research resulting in scientific publications authored by researchers based in large firms has declined since 1980, across industries, in both the US and Europe. One striking illustration of this pattern is that while 41% of R&D 100 awards went to Fortune 500 firms in 1971, only 6% went to those firms in 2006 (Block and Keller, 2009).

Nonetheless, the frequency with which large firms' patents cite scientific research has not declined. Arora, Belenzon, and Sheer (2021) suggests that these patterns are driven by the tradeoff between the use of science in own-inventions and the spillovers of inventions to rivals. As firms become more specialized and shrink their scope, they rely more on purchasing inventions created by small firms or on drawing from university science. Even when large firms are not doing their own pure science, their reliance on science can be traced through patent citation patterns. The majority of patents either cite scientific journals themselves or are indirectly connected to scientific publications through the network of patent-to-patent citations (Ahmadpoor and Jones, 2017). This is all to say that an important share of innovative activity is being performed by researchers motivated by the norms and institutions, in many cases non-market, of science. To understand innovation at large, we thus need to understand what affects the rate and direction of innovation produced by these researchers.

2.1 What drives the rate and direction of scientific research?

What motivates scientists producing science outside of private firms? How do they choose which topics to study? Sociologist Robert K. Merton[29] was perhaps the first to convincingly articulate a description of scientists being motivated by a desire to establish priority of discovery, as well as the key role of disclosure in communicating new discoveries and garnering recognition by other scientists for making the discovery first (Merton, 1957, 1968, 1969). In Merton's framework, publication of scientific results plays two key roles: the disclosure required to publish establishes priority, and citations to publications provide one, albeit imperfect, metric of how important a given scientific contribution is. While other rewards such as naming of scientific discoveries (e.g. Planck's constant) and occasional prizes carrying either prestige or monetary rewards or both (e.g. the Nobel Prize) may also accrue to scientists making new discoveries, in Merton's framework the importance of being first and thus establishing priority is the underlying incentive affecting behavior.

Stephan (2010) notes that science is sometimes described as a winner-take-all contest, meaning there are no rewards for being second or third even when competing research teams solve problems near-simultaneously (Bikard, 2020). Winning a priority race may benefit the victor over time by permitting easier access to resources including funding, graduate students, and reputation. In this way, one may worry that the "rich get richer" as those who publish early acquire the resources that lead to scientific production in the future. Merton (1968), drawing on scientist interviews conducted by Harriet Zuckerman, refers to this hypothesis as the "Matthew Effect."[30] Dasgupta (1989) and Dasgupta and David (1987) argue this winner-take-all feature may arise due to the difficulty of monitoring scientific effort. Hill and Stein (2020b) investigates the extent to which a winner-take-all model looks empirically relevant in the scientific field of structural biology. In surveys, scientists in this field estimate that being "scooped" by another team of scientists in a priority race to solve the same problem would face a 59 percent penalty in citations compared to the hypothetical winner. However, in a carefully constructed empirical test comparing winners and losers of priority races – who appear ex ante similar on observable factors – Hill and Stein estimate that scooped teams in fact face only a 20 percent penalty in citations. Taken together, these facts both reject a winner-take-all model in this

[29]Notably for economists, Robert K. Merton was the father of Robert C. Merton, the Nobel Prize-winning economist.

[30]Merton has frequently noted, e.g. see footnote 2 in Merton (1988), that Merton (1968) should have been co-authored with Zuckerman.

field of science, and suggest that scientists may overestimate the costs of being scooped.[31]

Bobtcheff, Bolte, and Mariotti (2016) theoretically investigates how the importance of priority affects research investments. The authors analyze a general model of priority races, conceptualizing the reward structure as winner-take-all, and explicitly incorporate the tension between letting a project mature for longer (improving its quality) versus patenting or publishing to establish priority. Hill and Stein (2020a) presents a related model which allows for endogenous entry across projects that vary in their ex-ante potential. High potential projects are more attractive because they offer higher payoffs, so researchers invest more in trying to enter those projects. This makes high potential projects more competitive, which in turn leads scientists to prematurely publish their findings. Hence, in this framework high-potential projects – that is, projects tackling research questions the scientific community has deemed most important – are the projects that will be executed with the lowest quality. Hill and Stein then present a series of empirical tests, again leveraging data from the field of structural biology, that provide evidence consistent with this model.

Priority is a fundamental feature of science, but so is autonomy. As articulated by Dasgupta and David (1994) and Merton (1973), scientific researchers tend to be offered substantial freedom in choosing projects. Stern (2004) takes this characterization of scientific research as a starting point for testing whether scientists have a "taste" for science, in the sense that they are willing to accept a lower wage in exchange for choosing their own research projects and participating in scientific communication. He analyzes data on job offers made to postdoctoral students in biology, and in an individual fixed effects framework documents evidence of a positive compensating differential. Plainly stated, scientists seem willing to pay to do science.[32]

This finding forms the starting point for the theoretical model of Aghion, Dewatripont, and Stein (2008), which seeks to clarify the relative advantages and disadvantages of different stages of scientific research (early-stage basic research versus later-stage applied research) being conducted at universities versus in the private sector. In the model, academia and private sector research are separated by differences in creative control (more prevalent in academia) and focus (more prevalent in private firms). If scientists value creative control more than focus, they will have to be paid a wage premium to give it up, as suggested by the evidence in Stern (2004).

The Aghion, Dewatripont, and Stein (2008) model points to tradeoffs in research funding faced by a social planner. Academia may be a cheaper input to innovation production, given the compensating differential. The incentives of academics, however, may not align with those of the social planner. Scientists may work on projects they find personally interesting or prestige-enhancing, but which have little economic value. Giving control rights to firms can direct scientists to work on more commercially valuable projects, but is more expensive because scientists must be paid a wage premium. In the model, the resolution of the tradeoff depends on how close to the point of commercialization the particular research stage is. The model also suggests that stronger property rights — often cited as a solution to underinvestment in innovation effort due to knowledge spillovers — may not be optimal. Murray, Aghion, Dewatripont, Kolev, and Stern (2016) documents empirical evidence consistent with the predictions of the Aghion, Dewatripont, and Stein (2008) model, in the context of research on genetically engineered mice.

Given the apparent value scientists place on research freedom, it is perhaps unsurprising that it appears to be challenging to use policy levers to change the direction of scientific research effort. Goolsbee (1998)

[31] Of course, the question of how scientists agree on what has been proven or shown is itself interesting. Latour and Woolgar (1986) describes how a Nobel-winning laboratory at the Salk Institute decides how internal or external research "counts" as a new result instead of an error or anomaly. Shwed and Bearman (2010) uses the network of citations over time to investigate empirically how scientific groups with contradictory results find consensus. Uzzi, Mukherjee, Stringer, and Jones (2013) investigates citations in the universe of Web of Science, finding that highly-cited work tends to be slightly atypical in the combinations of knowledge it draws on, but not very unusual. They argue that this is at least partly due to the fact that highly unusual new knowledge is very difficult for outsiders to interpret.

[32] Indeed, as the pre-economic literature on innovation made clear, innovators have many motivations beyond pure financial reward. For instance, when Wikipedia was blocked in mainland China, Chinese-language contributions to Wikipedia from Taiwan and Hong Kong fell substantially (Zhang and Zhu, 2011). Even though no payment was made for writing Wikipedia articles, social recognition among a large group of readers was an important inducement.

argues that increases in defense R&D cause substantial increases in the wages of scientists working in targeted sectors, with up to half the increased spending going to wages of those already working in the field. Myers (2020) analyzes US National Institutes of Health (NIH) Requests for Applications, which attempt to solicit research on particular topics in exchange for higher expected funding. He estimates that shifting one NIH applicant toward a project which is a relatively small distance (in science space) away from their prior research requires millions of dollars in higher funding.

While there is relatively limited research on precisely how scientists choose projects, Azoulay, Fons-Rosen, and Graff Zivin (2019) argues – based on an analysis of surprise deaths of star scientists – that scientists are hesitant to do work similar to that of prominent senior researchers (while the senior researchers are alive), at least in part because these powerful "stars" act as a deterrent on research shifting the direction of their field. To the extent that it is difficult to get scientists to change which project they work on, society may still benefit if the high-value projects chosen by scientists in one field are more likely than not to have spillovers to other areas. Such indirect spillovers appear to be empirically relevant: for example, Sampat (2015) estimates "serendipitous" spillovers, documenting that half of drug patents citing NIH-funded research are in a different disease area than that intended by the original grant.[33]

2.2 Knowledge production: The burden of knowledge hypothesis

Our focus in the previous subsection was on science as a static institution. However, the structure of knowledge production in science has been substantially changing over time, and as science evolves presumably so does optimal science policy (Jones, 2011). Jones (2009) presents two striking facts about changes in scientific knowledge production over time. First, the age at which a scientist produces her "first" innovation has increased steadily, from just over 30.5 in 1985 to around 31.4 in 1999. Second, the average number of inventors named on a patent has risen, from 1.7 in 1975 to 2.3 in 1999. Given the increasing age of scientists and the increasing size of teams, Jones asks whether these trends reflect incentives unrelated to productivity, or whether these trends are instead the result of an increasing "burden of knowledge" necessary to make scientific advances.

His argument starts with the observation that innovators are not born, but made, and that before inventing they must receive enough education to reach the frontier of knowledge. As the total stock of knowledge grows over time, reaching this frontier requires increasing quantities of effort. He highlights two potential effort responses: first, learning more by increasing the time spent in training; and second, specializing more and working more in teams.

Subsequent work has supported this burden of knowledge model, and has added additional granularity. Wuchty, Jones, and Uzzi (2007) leverages data on the near-universe of academic papers and US patents to show that work by teams is increasingly dominating — in frequency and quality — work by solo authors. Jones (2010) uses the World Wars as an exogenous disruptor of many scientists' education to show that the rising age of achievement for scientists can be attributed to the rising educational burden, rather than other factors such as changing demographics. Jones and Weinberg (2011) analyzes Nobel laureates in physics, chemistry and medicine, contrasting the former with the latter two. While in all three fields discoveries in early life have declined in frequency over time, this trajectory is less marked in physics. Jones and Weinberg argue this cross-field difference can be attributed to the introduction of the field of quantum mechanics in physics, which overturned much of the prior literature, thus reducing the burden of knowledge for physicists. Agrawal, Goldfarb, and Teodoridis (2016) finds specialization and team size increase in mathematical fields whose frontier expanded significantly with the arrival of Soviet mathematics in the West after the fall of the Iron Curtain. Brendel and Schweitzer (2019) also documents evidence consistent with the burden of knowledge hypothesis in the field of mathematics, and argues that researchers' response is consistent with a division of labor. Bloom, Jones, Van Reenen, and Webb (2020) shows evidence of declining productivity per

[33] As discussed below, Azoulay, Graff Zivin, Li, and Sampat (2019) documents a similar result.

researcher in semiconductors, agriculture, and biotechnology.[34]

Jones (2011) articulates an argument about why this set of findings motivates a new set of policy challenges. First, entry into science may be reduced by, or selection may be skewed by, the burden placed on young academics caused by the expanding knowledge frontier. Second, as research becomes more intricate and specialized, the burden of knowledge on gatekeepers of quality — patent examiners, grant reviewers, and so on — also increases. Finally, if policy and evaluation methods favor more specialized, narrow research ostensibly performed by single researchers, the economy will systematically disincentivize broader and more team-based research.

Expanded team research shows particular promise in mitigating the burden of knowledge, but also raises new questions around incentives and organization. For example, even with repeated interaction, teams involve fundamental inefficiency due to moral hazard with credit sharing (Che and Yoo, 2001). Attempts to reduce this moral hazard by rewarding team effort in and of itself can cause inefficient team rather than individual production when the summed credit given to team members exceeds unity (Valiki, Teodoridis, and Bikard, 2020). Large teams appear less likely to pursue highly-disruptive research, and hence both small and large teams may play important complementary roles in novel ideas and development of those ideas (Wu, Wang, and Evans, 2019). It remains an open question to separate how much of the increase in specialization and team production among innovators is due to a real increase in the difficulty of reaching the frontier as opposed to inefficient changes in scientific funding, credit allocation, and researcher training.

One final caveat to the burden of knowledge is that it is of course difficult to identify in the present important general purpose technologies in the future. Jacob Schmookler, writing in 1954 about a decline in inventiveness in that period, considered two new technologies which might permit growth to continue when the "unused fund of technical knowledge is drained" (Schmookler, 1954). His technologies? Atomic research and space travel. Despite the transistor having just been invented, the importance of the computer was not mentioned, let alone the long-run medical consequences of the 1953 discovery of the double helix.

2.3 How should science be funded?

One of the most important ways in which the public sector supports innovation is through the direct funding of research. Figure 2 shows that federally funded R&D used to be the primary component of total US R&D spending, but that since about 1980 the private sector has grown to account for a higher share of US R&D spending. In 2015, the US federal government accounted for around \$120 billion – around a quarter – of total US R&D by source of funds. This federal research support involved fifteen federal government departments and a dozen other agencies (National Science Foundation (2018), Table 4-15), but by far the largest two supporters were the Department of Defense (DOD) and the Department of Health and Human Services (HHS, which includes the National Institutes of Health or NIH).

Federal funding consists of two distinct activities – conducting research directly and supporting external research. The former involves, for example, federally funded research and development centers or FFRDCs such as Los Alamos National Laboratory, while the latter includes, for instance, NIH grants to university researchers. Figure 3 tabulates data on R&D spending and funding by sector, which illustrates one important asymmetry: the federal government funds more external research than it spends internally, whereas higher educational institutions spend more than they directly fund.

The natural question that arises with publicly funded research is the social value of this spending. How different is the set of innovations society has access to because of public research support? For reasons we discuss more below, systematic empirical analyses of these types of questions are challenging. The available evidence from qualitative case studies suggests that many or most scientific discoveries can trace their roots to

[34] That the amount of invention, particularly after adjusting for increased R&D, is falling is an idea with a long history. Griliches (1989) notes "the much more rapid rate of growth in national R&D expenditures than in total patenting and the implicit suggestion of diminishing returns" during the 1970s and 1980s. In the early 1950s, Stafford (1952) notes that patents per capita, and the number of new patent classes, had been falling since World War I.

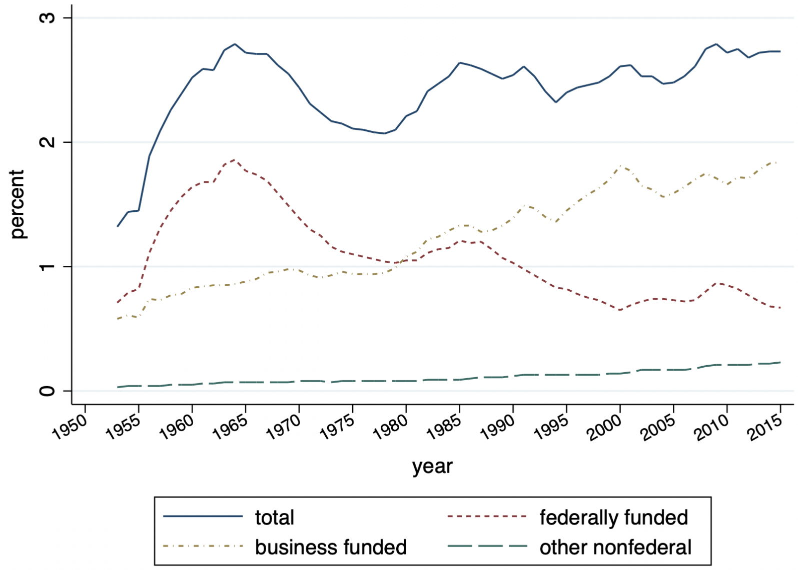

Figure 2: US research and development as a share of GDP, by source of funds: 1953-2015

Notes: This figure shows US research and development (R&D) spending – total and by source – as a share of GDP from 1953 to 2015. Source: Appendix Table 4-1 of National Science Foundation (2018).

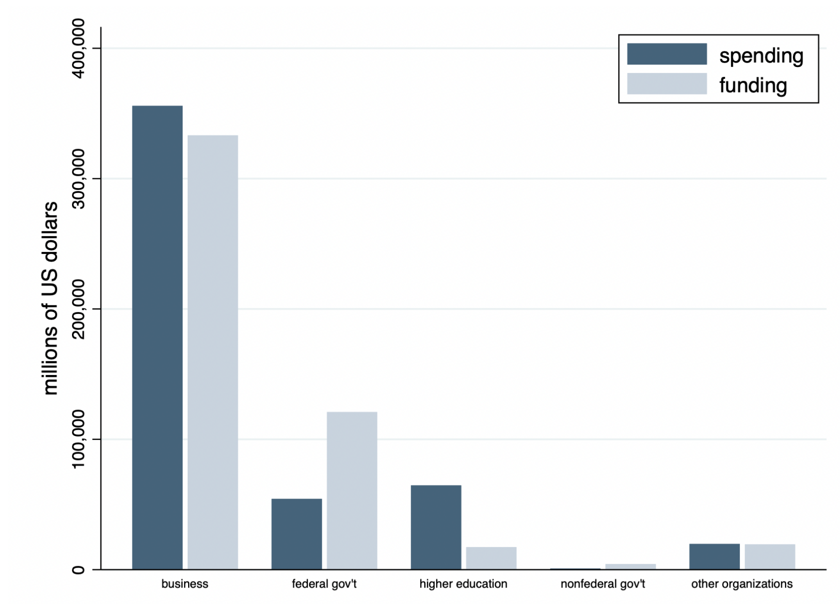

Figure 3: Spending on and funding of US R&D by sector: 2008-2015

Notes: This figure shows spending on and funding of US R&D by sector from 2008 to 2015. Source: Table 4-1 of National Science Foundation (2018).

both public research support and investment by private firms. For example, Chakravarthy, Cotter, DiMasi, Milne, and Wendel (2016) attempts to trace the history of individual drugs identified by a survey of physicians as the most transformative over the past twenty-five years, and argue that only four of the individual drugs appear to have been researched and developed solely by one sector.

These types of qualitative case studies are invaluable in providing textured examples. However, given the high level of public spending on research support, a central question of policy interest is what – quantitatively – the return is to incremental dollars of public research funding. Consider as an example the so-called war on cancer. In 1971, then-US President Richard Nixon signed the National Cancer Act of 1971. This act gave the National Cancer Institute (NCI, a branch of the NIH) special budgetary authority intended to allow the NCI to more flexibly address the disease which at the time was the second leading cause of death in the US. At the same time, Nixon requested an additional appropriation of $100 million for cancer research. Was this war on cancer "worth it?" Several decades later, many have declared the war on cancer a failure. However, rigorously assessing this question – or, more broadly, estimating the return to public funding of research – requires addressing several challenges.

A first challenge is that we often lack a paper trail for measuring the impacts of research investments. Say that the war on cancer funded investments into cell signaling that formed the basis for the development of Gleevec, a leukemia treatment; what data would allow us to trace out that connection? In recent years, researchers have made progress in several promising directions on this measurement challenge. Sampat and Lichtenberg (2011) was a pioneering paper in this area, developing both direct and indirect measures of whether drugs approved by the US Food and Drug Administration (FDA) built on publicly supported research. In terms of direct support, the authors measure whether the patents linked to individual FDA-approved drugs (as recorded in the FDA *Orange Book*)[35] are assigned to government agencies as well as whether they disclose a so-called government interest statement.[36] In terms of indirect support, the authors measure whether drug patents cite earlier patents assigned to government agencies or patents disclosing government interest statements, or cite published biomedical research articles that acknowledge public research support. Based on these data, Sampat and Lichtenberg estimate that while only about 9% of new drugs approved by the FDA between 1988-2005 directly benefited from public research support, nearly half (47.8%) benefited indirectly.[37] Constructing these types of linkages between public research support and new inventions is important partly because public investments tend to generate benefits with long and variable lags, and which are very diffuse and hard to predict in advance. Analogous to our discussion in Section 1.3.1, in the absence of direct linkages it would be difficult to know where to "look" for the set of technologies which benefited from public research support.[38] Perhaps the best evidence on this point comes from Azoulay, Graff Zivin, Li, and Sampat (2019) Table 9, which documents that NIH funding is as or more likely to generate private-sector patenting in other scientific areas as it is in the scientific area targeted by the grant.

A second challenge is that in many contexts, we would expect public research funding to be targeted at areas that are scientifically promising. This clarifies the need for a control group, particularly because any shifts over time in the scientific potential of a given area would be expected to affect incentives for both *private* research investments and public research investments. The potential interaction of private and public research investments also raises the issue of crowd-out: if a dollar of public research crowds out a dollar of private research, then public research investments could in theory have no real effect on total research spending, much less on health outcomes or overall welfare (David, Hall, and Toole, 2000).

The key advances on these challenges in recent years have been made in the context of understanding the impacts of NIH support for biomedical research in the US. Jacob and Lefgren (2011) analyzes data on all applications (unsuccessful as well as successful) to the NIH from 1980 to 2000 for standard research grants (known as R01 grants). Applications are assigned priority scores based on independent scientific

[35] Patents recorded in the FDA *Orange Book* are removed when they expire, so constructing a complete set of patents from the *Orange Book* requires reconstructing a list from each annual version of the publication. One of us (Williams) digitized the historical *Orange Book* patent and exclusivity tables for years 1985-2016 (no *Orange Book* was published in 1986), based on PDF versions obtained via a Freedom of Information Act (FOIA) request; those data are available here: https://www.nber.org/research/data/orange-book-patent-and-exclusivity-data-1985-2016.

[36] Section 202(c) of the US Bayh-Dole Act requires a statement in patent documents stating the existence of federal funding where relevant, a requirement referred to as a government interest statement. See Rai and Sampat (2012) and de Rassenfosse, Jaffe, and Raiteri (2019) for more detailed discussions of government interest statements, and Durvasula, Ouellette, and Williams (2021) and Long (2019) for recent analyses of government interest statements for pharmaceuticals.

[37] Looking even more broadly, Cleary, Beierlein, Khanuja, McNamee, and Ledley (2018) finds that every one of the 210 new drugs approved by the FDA between 2010 and 2016 had publicly-funded prior research on the entity itself or on its molecular target.

[38] The main alternative to these linkages is to specify a lag structure and test how funding for a given disease affects outcomes for that disease, as in Blume-Kohout (2012), Manton, Gu, Lowrimore, Ullian, and Tolley (2009), and Toole (2012).

reviews, and the authors document a strong nonlinear relationship between a proposal's priority score and the likelihood that an application is funded. In essence, this is a regression discontinuity approach, where grant applications just above and just below the priority score cutoff should be similar ex ante but differ ex post in their probability of receiving funding. Combining these data and design, Jacob and Lefgren estimate that receiving an NIH R01 grant (roughly $1.7 million) leads to one additional publication over the subsequent five years, about a 7% increase.

Azoulay, Graff Zivin, Li, and Sampat (2019) builds on the measurement approach developed by Sampat and Lichtenberg (2011) and the empirical approach pioneered by Jacob and Lefgren (2011) to estimate the impact of NIH funding on private-sector patenting by biopharmaceutical firms. Like Jacob and Lefgren (2011), their empirical approach leverages application-level data (measuring both unsuccessful and successful grant applications), but Azoulay, Graff Zivin, Li, and Sampat (2019) go beyond the regression discontinuity-type variation and construct a comparison based on the structure of the scientific review process at the NIH. The NIH is comprised of 21 institutes/centers (such as the National Cancer Institute), each of which receives congressional appropriations. However, scientific evaluation of applications occurs in so-called study sections. Study sections assign raw scores which are then normalized to ranks, and institutes/centers make funding decisions by comparing ranks across study sections up to the payline (the same discontinuous threshold analyzed by Jacob and Lefgren). This structure implies that grant applications judged by study sections to be very similar in quality can experience different funding outcomes. Linking grants to papers and citing patents, their baseline estimates suggest that $10 million in NIH funding leads to 2.7 additional patents.

Looking forward, we would highlight several areas as particularly promising for future work. First, much of the recent progress on both measurement and developing credible empirical approaches for estimating the return to public research spending has been in the health space, specifically leveraging data from the US NIH. But as noted above, even just within the US, federal research support involves fifteen federal government departments and a dozen other agencies. A good example is the work of Moretti, Steinwender, and Van Reenen (2019), which analyzes how defense-related research spending affects private research spending. In the longer run, Gross and Sampat (2020) uses World War II scientific funding to examine long-run shifts in the rate, direction, and location of US technological development. Taking a global perspective, public funding may be particularly important in more funding-constrained environments. Ganguli (2017) examines Soros grants following the dissolution of the USSR, using a grant eligibility cutoff to establish causality. These grants double publications for recipient scientists and make it much more likely that they remain working in scientific fields.

Second, in recent years there has been increasing interest in how this funding should be structured. One specific example has been the debate over whether public research support should be granted to people or to projects. Traditional public research grant mechanisms such as NIH R01 grants are project-specific grants, providing 3-5 years of funding intended to support researchers accomplishing a set of specific aims on a particular project. In contrast, programs like the Howard Hughes Medical Institute (HHMI) Investigator program provide funding over a longer period – five years, with an expectation of renewal – to a person who has flexibility over how to use the funds over time.[39] Theoretical work such as Manso (2011) articulates reasons why the latter type of contracts may differ in what type of research they incentivize – in particular, NIH R01-style contracts may motivate incremental refinements on past discoveries whereas HHMI-style contracts may provide researchers with more freedom to pursue radical and untested ideas. Azoulay, Graff Zivin, and Manso (2011) provides empirical evidence which is consistent with Manso's theoretical model. Ottaviani (2020) evaluates common grantmaking rules at the project level which are meant to limit political meddling or self-promotion by reviewers of work in their own field. Theoretically, there are a number of perverse

[39] Grants from the Rockefeller Foundation under Warren Weaver in the mid-20th century also emphasized funding the scientist rather than the project (Barany, 2019).

incentives that arise when fields differ in how "noisy" their applications are.[40] Optimal research contracting seems ripe for both theoretical and empirical work.[41]

Third, beyond interest in how funding should be structured, there has also been increasing interest in how funding applications should be evaluated. Institutions such as the US NIH rely heavily on peer review by field-specific experts. Li (2017) empirically analyzes the trade-off between such experts having an informational advantage in distinguishing more and less promising applications on one hand, and the fact that these individuals may also be most likely to have personal preference that impact – and potentially bias – their objectivity. She documents that evaluators are both better informed and more biased about the quality of projects closer to their own work, and estimates that on net the benefits of expertise weakly dominate the costs of bias. In a related paper, Li and Agha (2015) documents that higher peer-review scores are associated with better research outcomes, controlling for previous accomplishments of the applicants.[42]

While this suggests that peer review generates information about the promise of individual grant applications that may not otherwise be available, this work cannot speak to the relative merits of peer review compared to other allocation mechanisms. Peer review may disadvantage early career scholars and scholars from diverse backgrounds who are less likely to have access to grant writing support. Peer review may also do less well in evaluating the work of multidisciplinary teams (Jones, 2011). Moreover, peer review also imposes large costs on both applicants and reviewers. Could alternatives to peer review reduce bias and/or achieve the same outcomes at lower costs? Institutions such as the New Zealand Health Research Council and the Volkswagen Foundation have been exploring lottery-based grant allocation mechanisms as one alternative (Adam, 2019; Chawla, 2020), although ideally one could of course do better than a randomized allocation. Beyond the probability of winning a grant, evaluation delays and grant preparation time may distort researcher project choices. A funding program for Covid-19 research called Fast Grants, which promised to evaluate applications and distribute money within two weeks, found in a follow-up survey that 78% of their recipients would work on different projects if they were not constrained by funding, and that 44% would work on projects where there is less existing scientific consensus (Collison, Cowen, and Hsu, 2021). Further studies of the costs and benefits of grant-based, peer-reviewed scientific funding should prove valuable in helping us understand the merits and limitations of the status quo.

Fourth, dating back at least several decades there has been tremendous academic and policy interest in how best to encourage the development of innovations out of universities. A key policy in this area is the 1980 Bayh-Dole Act. Prior to Bayh-Dole, US federal agencies had inconsistent policies on their books about whether recipients of federal research grants could take title to inventions that came from federally funded research. Bayh-Dole sought to change that by allowing contractors such as universities receiving federal grants to obtain patents on inventions created by that research.[43] In Europe, the analogous policy change was the elimination of the so-called professor's privilege: under the professor's privilege, university researchers retained blanket rights to their inventions, and reforms generally shifted rights back to universities. Lach

[40]Optimal selection of applicants is a longstanding problem. The philosopher C.S. Peirce gave a marginal cost/marginal benefit analysis of this question in the mid-1800s (Peirce, 1879). More recently, the question of how much leeway to give biased screeners in selecting projects has been an active area of applied microeconomic theory research (e.g., Frankel (2021)).

[41]See Azoulay and Li (forthcoming) for a detailed review particularly of the empirical literature in this area. Liu et al. (2018) finds that peak scientific impact, like peak creative impact in the arts, tends to be concentrated in a very short period of a given scientist's career, that this period is hard to predict, and that productivity as measured by pure research output is no higher during this "hot streak." These facts collectively suggest that it may be difficult for funders to select high-value researchers purely based on past output. On the theoretical side, a number of recent papers on strategic experimentation apply to the question of when and why a funder may wish to force grantees to disclose successes or failures (e.g., Bonatti and Hörner (2011), Curello (2021), Halac, Kartik, and Liu (2017), and Keller, Rady, and Cripps (2005)). There has been very little empirical application of this theory to innovation policy thus far.

[42]Note that the NIH tends to use highly specialized panels. Gush, Jaffe, Larsen, and Laws (2015) looks at a New Zealand public funder and finds no link between panel scores of research applications and subsequent outcomes. Newham and Midjord (2018) considers a shift from sequential to simultaneous voting on project funding, finding substantial herding behavior in expert evaluations.

[43]See Hemel and Ouellette (2017) as well as Eisenberg (1996) and Mowery, Nelson, Sampat, and Ziedonis (2015) for insightful discussions.

and Schankerman (2008) is a classic paper in this literature, documenting evidence that in the US context university licensing income was increasing in researcher's royalty share; however, Ouellette and Tutt (2020) argues this result is driven by errors in the coding of university policies, and that once corrected there is no evidence that royalties impact outcomes at US universities.[44] Particularly given this exchange, more recent attention has focused on what can be learned from the European experience. For example, Hvide and Jones (2018) analyzes a 2003 policy change in Norway which shifted income from businesses and patents towards universities and away from individual researchers at those universities, which they argue led to a dramatic decline in start-ups and patenting among university researchers. Given the active debates around how universities can best encourage and support innovative activity, this area is ripe for additional work.

Finally, just as public roads and bridges support private economic activity (Asher and Novosad, 2020; Brooks and Donovan, 2020), large-scale scientific infrastructure investments may support other scientific investment. In addition to funding research applications or scientific education, governments also fund large, shared infrastructure such as the early internet, particle colliders like the CERN Large Hadron Collider (LHC), and biological repositories. This infrastructure affects the geography of science: for example, Helmers and Overman (2017) finds that a very expensive synchrotron diamond light source in the UK attracted scientific activity, including related research which did not directly use the synchrotron. That said, measuring the relative efficiency of infrastructure investments versus direct funding in the aggregate is challenging. In the long-run, physical capital appears more replaceable than human capital. Waldinger (2016) finds no long-run relationship between physical capital destruction during World War II and scientific output, whereas the dismissal of elite scientists from their posts before the war led to permanent declines in science in a given city.

Summing up, scientific norms and institutions play a complementary role to market incentives for the creation of new ideas. The efficiency of science depends on concerns like priority and scientific taste that may not perfectly align with social goals, may be getting more difficult necessitating larger teams and increased specialization, and depends in subtle ways on the details of funding mechanisms and researcher autonomy. Understanding and optimizing these non-market policies is critical for a well-functioning innovation system.

3 Theory and evidence on market-based innovation policies

Let us now turn away from science toward the role of market incentives on innovation. Dating back at least to the work of Schmookler (1966), economists have focused on the idea that "invention is largely an economic activity which, like other economic activities, is pursued for gain." Yet, looking back at the 1962 NBER volume, policies that affect inventor behavior by changing prices or restructuring incentives are notably absent (National Bureau of Economic Research, 1962). From today's perspective, both intuition and empirical evidence tell us that changing the costs and rewards facing firms can substantively change their research investments.

Examples abound from both history and the present day of inventors responding to incentives. Daniel Treadwell invented his screw machine in response to the embargo of 1807, and the engineer Richard Gatling, known from his farming inventions until the 1860s, shifted to military implements during the US Civil War when he invented his famous gun (Khan and Sokoloff, 1993). Textile inventions augmenting Indian cotton rather than US cotton became more common during the supply disruptions of the US Civil War (Hanlon, 2015). In a more modern case, Acemoglu and Linn (2004) documents compelling evidence that as the US population has aged that investments in drug development have shifted towards developing more drugs to treat diseases common in older age groups.

This idea that research investments respond to incentives is such an intuitive idea that it is important to clarify what view is *rejected* by the type of evidence documented by Acemoglu and Linn (2004). Qualitative

[44] Replication data for Ouellette and Tutt (2020) is available here: https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/TTGFW2.

historical accounts of important discoveries often characterize innovations as happening via fortunate accidents, such as Archimedes' Eureka moment.[45] If most scientific advances were to happen via lucky accidents, economics – and policy – would have little role to play. But economic theory clearly suggests incentives should drive research investments, and the available data support that idea.

When market rewards matter, we need to consider the design of public policies which shift the rate and direction of inventive activity. At a conceptual level, the government has two levers: reduce the costs of research, or increase the expected revenues from research. Costs can be reduced through tax credits, direct subsidies covering the costs of research, or labor market policies such as immigration changes that make it easier for firms to hire specialized workers with expertise in specific technological areas. Expected revenues can be increased by changes in antitrust policies or through awarding market power or payoffs directly via the patent system or prizes. In this section, we review some of the key evidence and open questions on these policy levers.

Note that our focus differs from many other reviews of the economics of innovation. In particular, we focus much more on policy levers and institutions than on pure supply and demand factors like market size, firm size, and appropriability. This focus is not because those factors are unimportant; rather, it reflects a shift in the focus of the economics literature toward policy evaluation in addition to laissez faire firm behavior. Cohen (2010) gives a detailed overview of the literature linking pure firm characteristics and demand factors with innovation.

3.1 Taxes and innovation

Taking as given the idea that research investments in a market economy may be too low due to market failures, tax-based subsidies are one natural response. As Hall (forthcoming) articulates, advantageous tax treatment of research expenditures is a market-oriented response which leaves the choice and pursuit of R&D investments with the private sector – in contrast with, e.g., direct public R&D spending, where the government tends to play a larger role in choosing which projects are funded.

Several aspects of the US tax code are thought to significantly affect R&D expenditures.[46] Since the mid-1950s firms have had the option of fully deducting R&D in the year it occurs, in contrast with expenditures on buildings or equipment which are deducted over a longer period, thus primarily reducing tax liabilities further in the future. This distinction means that the basic structure of the tax code subsidizes current R&D expenditures more than longer-term capital investments. In addition to this subsidy, many countries choose to provide additional tax incentives for R&D investments, such as R&D tax credits. In 2015, 28 of the 34 OECD countries as well as a number of non-OECD countries gave advantageous tax treatment to business R&D expenditures (Appelt, Bajgar, Criscuolo, and Galindo-Rueda, 2016). Figure 4 provides one characterization of the cross-country variation – as of 2017 – in tax support for business R&D, as measured by forgone revenues and refunds. Some countries like Finland have very few tax incentives, whereas in the Netherlands they comprise a very high share of total public R&D support.[47]

In the US, the federal R&D tax credit was first introduced in 1981 (US Department of the Treasury, 2016).[48] For many years the federal R&D tax credit was a temporary policy; it was extended 16 times between 1981 and 2015 before being made permanent as part of the Protecting Americans from Tax Hikes

[45] As Schaffer (1994) notes, these Eureka moments, from Archimedes to Kekulé's famous discovery of the structure of benzene, look far less serendipitous when examined ex-post by historians of thought rather than in the autobiographical tales of the scientist.

[46] See the discussion in Mansfield (1982).

[47] Hall and Van Reenen (2000) provides a clear and useful exposition of the many various dimensions on which the tax treatment of R&D can differ. These include how R&D is defined for purposes of a tax credit, carryback and carryforward provisions, and whether tax credits apply to levels of R&D or incremental changes compared to previous years.

[48] Though this credit was not introduced until the 1980s, R&D tax policy was the first recommendation in Bush (1945) "Science, the Endless Frontier": "Government should provide suitable incentives to industry to conduct research, (a) by clarification of present uncertainties in the Internal Revenue Code in regard to the deductibility of research and development expenditures as current charges against net income, and (b) by strengthening the patent system so as to eliminate uncertainties which now bear heavily on small industries and so as to prevent abuses which reflect discredit upon a basically sound system."

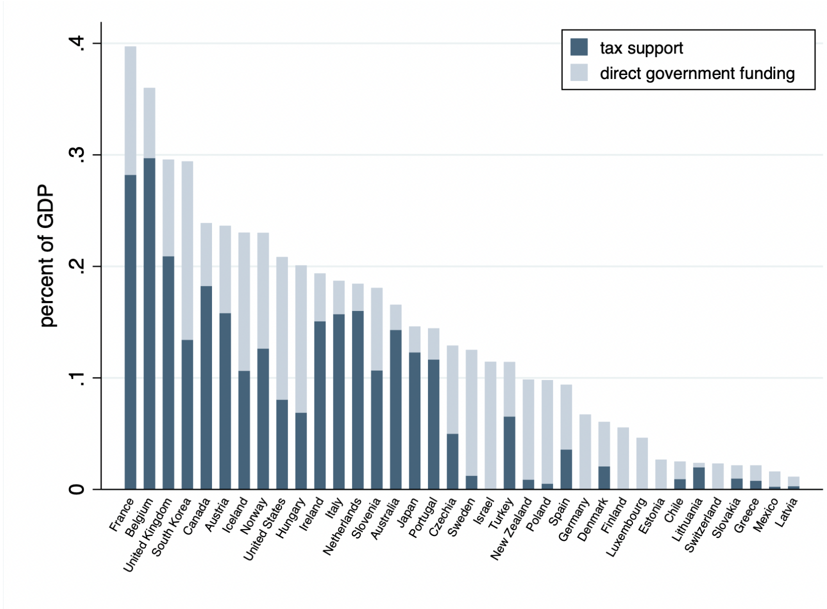

Figure 4: Direct government spending and government tax support for business R&D in 2017

Notes: Government tax support combines national and sub-national tax support for business R&D expenditure. Data on national tax support is not available in Israel. Data on sub-national tax support is not available in the US and Spain. Source: Organization for Economic Co-operation and Development (2020).

Act of 2015 (National Science Foundation, 2018). From a tax expenditure perspective, the federal R&D tax credit is estimated to cost the US federal government around $11 billion each year in foregone tax revenue (National Science Foundation, 2018). Minnesota was the first state to introduce a state-level R&D tax credit in 1982 (Miller and Richard, 2010). As illustrated in Figure 5 (which replicates Figure 1 from Wilson (2009)), in recent decades state-level R&D tax credits have become both more common and more generous.[49] Dating back at least to the work of Mansfield (1986b) and Hall (1993), economists have attempted to use such variation in R&D tax credit policies to try to shed light on the effectiveness of these policies.

Perhaps the most natural first question is whether R&D expenditures increase when the tax-adjusted user cost of R&D falls. A number of papers have estimated this relationship in country-year or state-year panel data. For example, Bloom, Griffith, and Van Reenen (2002) estimates the relationship between R&D spending and the tax-adjusted user cost of R&D in a panel of nine OECD countries from 1979-1997. More recent papers such as Dechezleprêtre, Einiö, Martin, Nguyen, and Van Reenen (2016) and Agrawal, Rosell, and Simcoe (2020) are able to leverage within-jurisdiction variation by using administrative tax records to analyze discontinuous firm size thresholds in the generosity of R&D tax credits. Looking across the available macro- and micro-estimates, Bloom, Van Reenen, and Williams (2019) argues that a reasonable summary of the estimated elasticities found in this literature is that a 10 percent fall in the tax price of R&D generates at least a 10 percent increase in R&D in the long-run.

Two types of concerns have generally arisen with such studies. First, because they have tended to use self-reported R&D expenditures as the outcome of interest, there has been a concern that measured R&D responses to tax changes may reflect re-labeling of existing expenditures as research expenditures rather than true changes in research investments. At a basic level, what expenditures should or should not be labeled as R&D is difficult to define, and self-reported R&D investment figures thus a priori seem likely to be subject to manipulation in response to incentives. Chen, Liu, Suárez Serrato, and Xu (2021) documents direct evidence of such re-labeling in China: a Chinese policy that awards substantial tax cuts to firms with R&D investment

[49] Wilson (2009) constructs the data for this figure from a variety of sources, the primary of which was online state corporate tax forms.

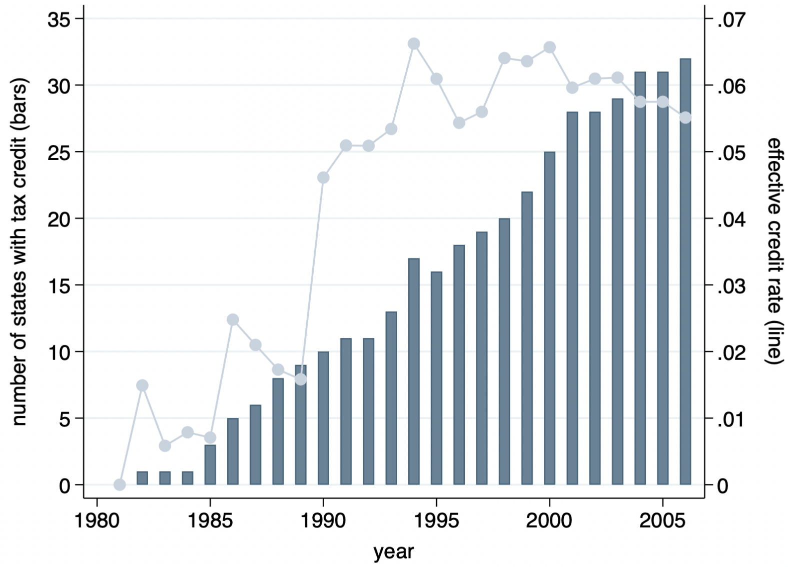

Figure 5: Number and average value of state R&D tax credits in the US, 1981-2006

Notes: This graph replicates Figure 1 from Wilson (2009).

over a threshold results in a significant increase in reported R&D, around a quarter of which is driven by firms relabeling expenses as R&D.

Dechezleprêtre, Einiö, Martin, Nguyen, and Van Reenen (2016) addresses this concern by looking not just at R&D expenditures but also looking directly at "real" outcomes in UK data including patenting and total factor productivity. They find that patenting, along the treated margin of relatively small firms, has an elasticity of 2.6 with respect to tax-adjusted user cost, and further that this R&D spills over to technologically-similar firms who do not directly receive the credit. In similar data, Pless (2021) investigates the relationship between R&D grants and tax credits. Small firms are more likely to be liquidity-constrained, hence direct grants are complementary to other R&D investments. For instance, small firms may use their grant to hire researchers who work with expensive capital equipment that is subsidized by the tax credit. On the other hand, larger firms tend to use the grants to pay for inframarginal investment, making the tax credit a substitute rather than a complement. Drawing on recent advances in public finance, mechanism design approaches to more precisely target R&D taxes and subsidies deserve further investigation (see Akcigit, Hanley, and Stantcheva (2019) for a structural example).

A separate concern is that country- or state-level R&D tax credit policies may not raise aggregate R&D but may rather simply incentivize the relocation of research investments to areas with more generous tax incentives. Wilson (2009) documents evidence consistent with a firm mobility response: lower R&D tax rates attract more R&D investment, but this appears to be offset by lower R&D spending in high tax rate states. Three more recent papers have tested for evidence of relocation at the level of individual scientists/inventors: Moretti and Wilson (2014) tests whether state-provided financial incentives for biotech companies raise the number of biotech scientists in a state; Akcigit, Baslandze, and Stantcheva (2016) analyzes inventor mobility in response to country-level top tax rates; and Moretti and Wilson (2017) analyzes mobility in response to state-level tax rates within the US (both personal and corporate income taxes). Taken together, these papers provide a compelling case that relocation is a relevant margin at least in some contexts. For example, Akcigit, Baslandze, and Stantcheva (2016) documents that after the collapse of the USSR, Russian inventors were far more likely to settle in countries with lower top marginal tax rates. From a research validity perspective, the work of Moretti and Wilson (2014) and Moretti and Wilson (2017) finds, comfortingly, that scientists/inventors seem to respond to tax incentives only when these policies apply to them (e.g. corporate tax changes have no effect on academic or government scientists).

Of course, this evidence of a relocation response does not imply that the aggregate effects of R&D tax credits are zero. But quantifying the aggregate effects of R&D tax credits given evidence of relocation responses is challenging. Analyzing the impact of corporate and personal taxes on innovation in the US over the course of the twentieth century, Akcigit, Grigsby, Nicholas, and Stantcheva (2018) argues that there is evidence for relocation but that the aggregate effects of tax incentives on innovation are nonetheless substantial. Also related is a pair of recent papers, Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019b) and Jones (2019), which analyze at a more conceptual level whether top income tax rates should be expected to have aggregate effects on R&D. In a similar vein, as best we are aware there is little work attempting to relate the magnitude of these tax policy responses to the magnitude of the gap between social and private returns to research investments, as estimated by Bloom, Schankerman, and Van Reenen (2013) and others. Looking forward, these questions are promising areas for future work.

A related topic which we also expect to be a promising area for future work is patent box policies. First introduced in Europe in the 1970s, patent boxes apply a lower tax rate to revenues linked to patents; the policy has since been implemented in many European countries, including France, Ireland, and the United Kingdom. While the policy is often described as a way of incentivizing R&D, in practice patent boxes raise the risk of inducing tax competition given that firms – particularly multinationals – often have flexibility in deciding where to book the taxable income generated by patented products. Consistent with this concern, Gaessler, Hall, and Harhoff (2018) and Alstadsæter, Barrios, Nicodeme, Skonieczna, and Vezzani (2018) provide evidence of high-value patents being relocated to markets with patent boxes, and fail to find evidence of either higher R&D expenditures in markets with patent boxes or aggregate increases in R&D investments.

At a conceptual level, Hall (forthcoming) provides an insightful comparison of R&D tax credits and patent boxes: R&D tax credits target inputs (R&D spending), whereas patent boxes target an output (profits) and (unlike R&D tax credits) only cover patentable innovations. Hall also points out that patent boxes effectively subsidize patent assertion – given that the income of firms that specialize in patent litigation/enforcement is patent income – and provide incentives to renew patents that would otherwise be abandoned. Bloom, Van Reenen, and Williams (2019) argues that both logic and the available empirical evidence suggest patent boxes are likely to be a harmful form of tax competition that should be discouraged, relative to well-designed R&D tax credit policies.

A final area we would highlight is the direct public funding of research at private firms. Governments routinely subsidize research investments at private firms, particularly small firms, yet until recently we had very little evidence on the efficiency of these subsidies. In a classic paper, Lerner (1999) provides an empirical analysis of the US Small Business Innovation Research (SBIR) program, arguing that program awardees grew significantly faster than matched controls. Howell (2017) provides a more recent analysis of the same program, leveraging administrative data in a regression discontinuity approach to document evidence that early-stage SBIR awards have large, positive effects on firms' probability of receiving venture capital support as well as on firm patenting and revenue. Bronzini and Iachini (2014), Einiö (2014) and Santoleri, Mina, Di Minin, and Martelli (2020) provide evidence on similar programs in Europe, as does Le and Jaffe (2017) for New Zealand. An important caveat to the benefit of direct funding comes from Bhattacharya (forthcoming): government agencies may be maximizing objectives other than overall social welfare. He finds in a structural auction model that Department of Defense SBIR contracts are inefficient, but that this may be in part because the DoD is minimizing acquisition costs rather than maximizing overall surplus including that of inventors. Several of these papers have undertaken extraordinary efforts to obtain access to administrative datasets that enable novel and compelling empirical approaches, which we hope will pave the way for additional work in this area.

3.2 Intellectual property rights

Formal intellectual property (IP) is a set of government policies which grant special rights to exclude others from producing, selling, licensing, or building on products, resulting in less competition than would be faced in a free market. Unlike taxes and subsidies, which affect the cost of doing R&D, IP affects both the cost of research by sometimes requiring inventors to license ideas they wish to build on, and the reward to research by limiting ex-post competition. Though IP is a broad area, we will focus our attention in this subsection on patents.[50]

3.2.1 Patents: A primer

Many of the key open questions about the patent system are easier to articulate against the background of some basic facts about the system.[51] To be concrete, we focus on the US Patent and Trademark Office (USPTO), but we bring in details about the international dimensions of patenting at several points.

When filing a new patent application, inventors must submit a written description of their invention which includes a discussion of so-called prior art – publicly available information relevant to the originality of the invention being patented (such as might be codified in previously filed patent applications or scientific publications). Patent applications also include a specific list of claims that the applicant seeks to assert intellectual property rights over. Applicants pay filing fees that vary with the type of patent application being submitted as well as whether the applicant qualifies as a small or micro entity (who are allowed to pay lower fees).[52] The full cost of filing a patent, inclusive of legal drafting, averages between $10,000 and $40,000 in the US, a nontrivial barrier particularly for startups and other small firms (Graham, Merges, Samuelson, and Sichelman, 2009).[53]

Once a patent application is submitted, it is assigned to a group of patent examiners with relevant technological expertise (an art unit), and then assigned to a specific examiner for review. The examiner is responsible for determining whether the patent application meets the standard for patentability – namely, whether it is patent-eligible (Title 35, US Code §101), novel (Title 35, US Code §102), nonobvious (Title 35, US Code §103), useful (Title 35, US Code §101 and §112), and whether the text of the application meets the requirements for disclosure and claim definiteness (Title 35, US Code §112). Examiners can choose to issue a first action allowance of the patent (that is, granting the patent based on the initially submitted application), but this is relatively rare. More often, the initial decision is a rejection, after which applicants have the option to submit a revised version of their application. For example, the examiner may reject specific claims and explain their rationale; because claims can be altered or eliminated in response, this step is important given that the number and content of patent claims is what determines the breadth or scope of a patent (Marco, Myers, Graham, D'Agostino, and Apple, 2015). The patent application review process often involves multiple rounds of rejection and revision. In practice, patent applications cannot be rejected by the USPTO, only abandoned by applicants (Lemley and Sampat, 2008).

In patent reform debates, attention often focuses on whether the USPTO is granting "too many" or "too

[50] Other forms of IP available in some jurisdictions include trade secrets, copyrights, trade dress, geographical indications, trademarks, design rights, and database rights.

[51] See Williams (2017) for more details.

[52] For more on the USPTO's fee structure, see https://www.uspto.gov/learning-and-resources/fees-and-payment/uspto-fee-schedule.

[53] Graham, Merges, Samuelson, and Sichelman (2009) derive this cost estimate from a survey of high-technology startup firms in the US, and it is of course not necessarily representative of costs faced by e.g. independent inventors. In addition to this literature that has estimated the total costs of filing, there is evidence that filing fees also help weed out low quality patents, reducing the burden on the patent office. Nicholas (2011) examines an 84% reduction in filing fees in the United Kingdom which occurred after the 1883 Patents Act, and finds that patenting increased substantially but that the number of important innovations created in that country did not change. de Rassenfosse and Jaffe (2018) considers a fee increase in 1982 in the US, and finds a reduction in the patent propensity of low-quality inventions, especially from firms with large portfolios.

few" patents.[54] These arguments are not new: see Machlup and Penrose (1950) for a discussion of the 19th century patent abolition movement, and Polanvyi (1944) for conceptual arguments against patents from the age of industrial research. Modern debates about patents in practice often draw on, e.g., the subjectivity of the novelty requirement – if the USPTO is using too low of a bar for novelty, it may grant too many patents. However, before we can empirically assess whether the patent grant rate is too high or too low, we need to measure what the patent grant rate is, and it turns out that measurement exercise is itself quite challenging.

Basic data on patent applications and granted patents by year are widely available: Figure 6 tabulates counts of USPTO applications and grants by year from 1963-2019. Until recently, data on unsuccessful patent applications were not made public, complicating measurement of what share of filed patent applications were granted patents. However, even in datasets where unsuccessful applications can be accounted for, the structure of the patent examination process means there is not a well-defined time window during which a final allowance or rejection decision will occur. Multiple rounds of rejection and revision common to most applications can unfold quickly or slowly depending on the USPTO's review lags and applicants' timeliness in resubmitting. Moreover, patent applications frequently produce new but closely related patent applications such as continuations or divisionals. These should arguably be incorporated into a measure of whether the ideas in the initial progenitor patent application were granted.

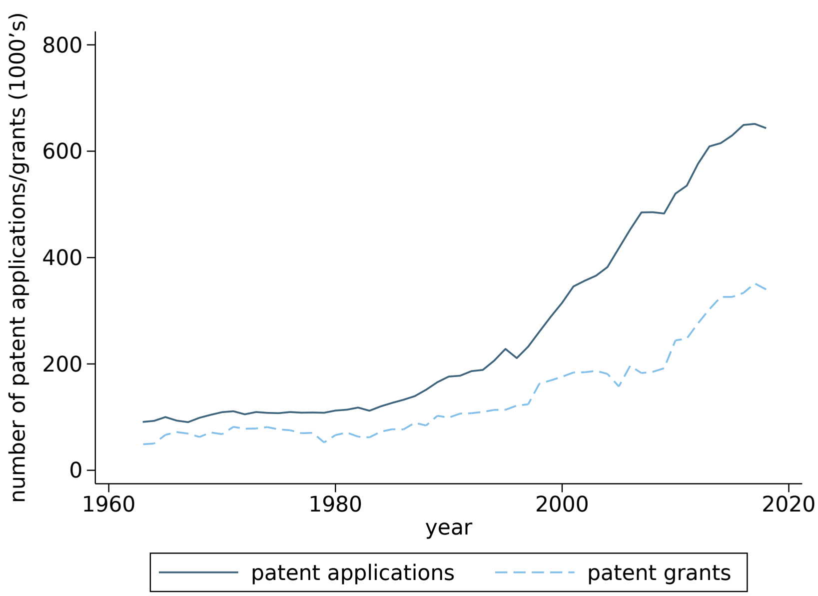

Figure 6: USPTO patent applications and grants by year, 1963-2019

Notes: These totals include utility, design, and plant patents as well as patent reissues. Source: USPTO Patent Technology Monitoring Team (2021), which tabulates data from USPTO submissions to World Intellectual Property Organization (WIPO), USPTO's Technology Assessment and Forecast (TAF) database, *Statistical Abstracts*, and *Annual Index Patents*.

Figure 7 replicates Figure 2 of Carley, Hegde, and Marco (2015), which uses non-publicly available USPTO administrative data, including all applications successful or not, to attempt to overcome these

[54] Note that this concern masks two conceptually distinct issues. First is that the USPTO may have an implicit standard that is either "too high" or "too low" by some normative judgement. Second is that the USPTO may be inconsistent in its application of its own standard. de Rassenfosse, Jaffe, and Webster (2016) provides one attempt at distinguishing between these two alternatives, estimating that the implicit threshold of the USPTO is modestly lower than that of peer patent offices in Europe and Japan, and that the fraction of patents in the US that are granted inconsistently is relatively low (around 20 percent).

challenges.[55] Their analysis sample is constructed from so-called progenitor patent applications (roughly, patent applications not derived from previous applications) filed with the USPTO from 1996 to 2005 and examined before mid-2013. The first action allowance rate line shows what share of applications receive a first-round acceptance. The progenitor allowance rate measures the share of patent applications that are directly allowed. The family allowance rates incorporate patent grants that accrue to continuations or divisionals derived from the original application.[56] These data confirm that initial allowance rates are relatively infrequent (consistently less than 20%). Progenitor allowance rates are much higher, giving a sense of the fact that many applications are granted patents only after at least one round of revision, and family allowance rates are higher still.

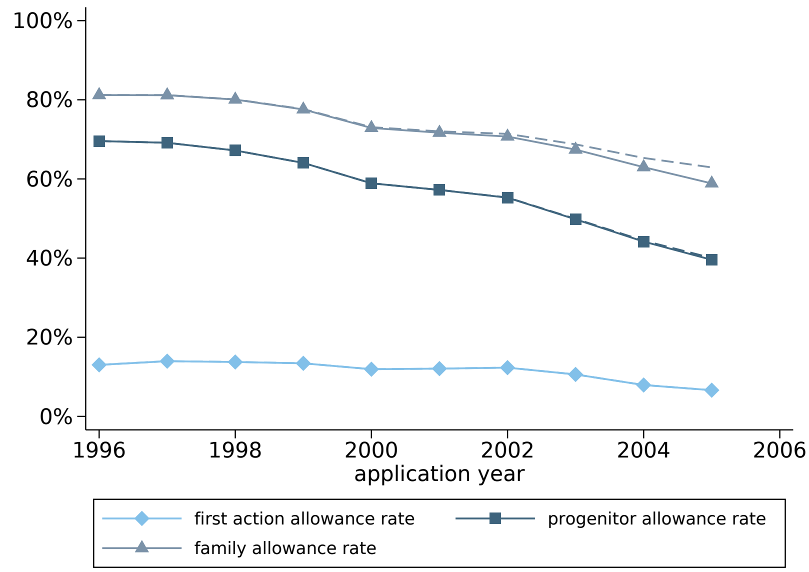

Figure 7: USPTO allowance rates for applications filed from 1996-2005 and examined before mid-2013

Notes: The first action allowance rate is the share of patent applications allowed at their first substantive examination. The progenitor allowance rate is the share of patent applications directly allowed. The family allowance rate is the share of progenitor applications that produce at least one granted patent, including continuations and divisionals. The dashed lines represent the highest possible allowance rates (i.e., the allowance rate assuming that every pending application as of June 30, 2013 is allowed). Source: Carley, Hegde, and Marco (2015).

The value distribution of granted patents is quite skewed, as shown by Pakes (1986). Pure counts of patents by firm or region can mislead if this skew is not accounted for. A variety of metrics have been proposed to measure the value of granted patents, including forward patent citations (Trajtenberg, 1990), patent renewals (Bessen, 2008; Pakes, 1986; Schankerman and Pakes, 1986), licensing revenue (Abrams, Akcigit, and Grennan, 2019), and excess stock market returns (Kogan, Papanikolaou, Seru, and Stoffman, 2017).[57] These metrics measure conceptually different quantities: for example, the Kogan, Papanikolaou, Seru, and Stoffman (2017) excess stock market returns measure focuses on very narrowly measuring the "surprise" component of the private returns to a patent that is realized at the time the patent is granted. In contrast, forward patent citations as analyzed by Trajtenberg (1990) and Bloom, Schankerman, and Van

[55] A separate issue is that because the text and claims of patents are revised during the examination process, grant decisions are not binary outcomes. Credit to Ben Roin for very insightfully pointing out in an unrecorded conference discussion that presumably any patent application could be granted a patent if the claims were sufficiently narrowed, but that at some point claims that are too narrow provide little economic value to applicants and those applications are then presumably abandoned.

[56] Importantly, note that "family" here refers to a different object than patent families in the international context, in which families refer to groups of patent applications filed in different jurisdictions which cover the same technology.

[57] See also Lanjouw and Schankerman (2004).

Reenen (2013) aim to measure spillovers, which more closely align with social value rather than private value. Figure 8 nonetheless – and perhaps surprisingly – shows a strong positive correlation between forward patent citations and the Kogan, Papanikolaou, Seru, and Stoffman (2017) excess stock market returns measure.

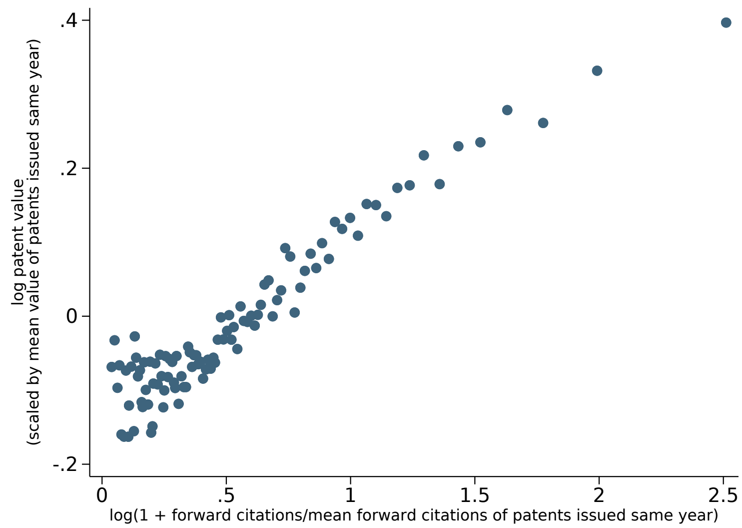

Figure 8: Forward citations and market value (1926-2010)

Notes: This figure plots the cross-sectional relationship between forward patent citations and the estimated excess stock return patent value estimate constructed by Kogan, Papanikolaou, Seru, and Stoffman (2017), which is defined only for granted US patents assigned and linked to publicly traded firms. Sources: Kogan, Papanikolaou, Seru, and Stoffman (2017) and USPTO administrative records.

USPTO-granted patent rights generally do not extend to foreign markets, so inventors wanting equivalent rights in other countries must file a patent application covering the same invention in that country or region (a set of patent applications covering the same invention is referred to as a patent family). Once patents are granted, owners must periodically pay maintenance fees in order to keep their patent active.[58] Conditional on a patent being active, owners have the authority to enforce their patent rights through litigation. Only 1-2% of patents are ever litigated (Bessen and Meurer, 2005; Chien, 2011; Lemley, 2001), although of course more are challenged that settle out of court. Both theoretical and empirical analyses have emphasized that more valuable patents are more often targets of legal action (Allison, Lemley, Moore, and Trunkey, 2004; Allison, Lemley, and Walker, 2009; Cremers, 2004; Lanjouw and Schankerman, 2001; Miller, 2013; Moore, 2005; US Patent and Trademark Office, 2015). Patents can be bought and sold, and owners have the option – although no requirement – to record ownership changes with the USPTO.[59] To the best of our knowledge, Serrano (2010) is the first to analyze the USPTO's data on recorded ownership changes. Figure 9 replicates Figure 1 from Serrano (2010), plotting the share of active patents with a recorded ownership change by age of the patent, separately for patents which have been previously traded versus not. The key takeaways from this graph are that previously traded patents are more likely to be traded again, and that for both groups the probability of ownership changes declines with patent age.

From a practical perspective, digital full-text versions of patent grants since 1976 are publicly available,[60] and in more recent years essentially each step of the application process is publicly documented in the PAIR (Patent Application Information Retrieval) data.[61] These types of detailed administrative data enable

[58]See Scotchmer (1999) on the optimality of the patent renewal system, both timing and fee amounts.

[59]Graham, Marco, and Myers (2018) notes that the extent and reasons why some changes in assignment are not recorded is an open question.

[60]See https://bulkdata.uspto.gov.

[61]See https://www.uspto.gov/learning-and-resources/electronic-data-products/patent-examination-research-dataset-public-pair.

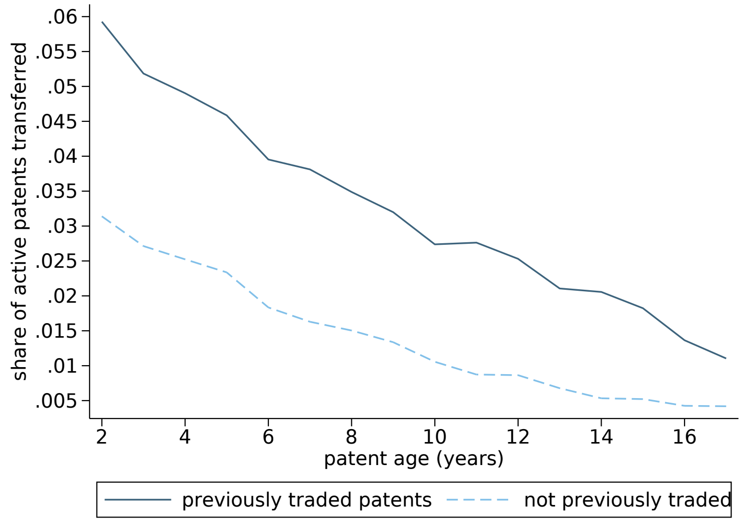

Figure 9: Share of active patents traded conditional on having been previously traded or not

Notes: This figure is similar to but not exactly the same as Figure 1 from Serrano (2010). Source: Marco, Myers, Graham, D'Agostino, and Apple (2015).

modern text analysis methods to, for example, compute text-based measures of patent quality as in Kelly, Papanikolaou, Seru, and Taddy (forthcoming). From these full-text data, data can be extracted on everything from citations (front-page or in-text) and inventor locations to assignees and claims.[62] Derived data such as the PatentsView data provide some pre-processed versions of these variables, but of course in many cases researchers may need to extract different versions of variables from the raw data files.[63]

3.2.2 Intellectual property: Theory

As we have seen, a patent grants an inventor the right to exclude others from economically exploiting an invention for a limited period of time, in exchange for disclosing how the invention functions.[64] Formal legal protection of ideas is old, with patents dating back to 15th century Venice (Comino, Galasso, and Graziano, 2020). That said, as Plant (1934) notes, patents are quite strange. We normally think of property rights as a solution to scarcity. Patents, however, take a nonrival resource called knowledge and create an artificial scarcity. Other formerly-common policies which create artificial scarcity, such as government-granted monopolies and charters, are now seen in a negative light. An understanding of the theoretical rationale for patents, both in and of themselves and relative to secrecy, is therefore essential.

Theoretically, patents can be broadly understood as providing four benefits: they allow inventors to appropriate more of the social value of their direct invention, reward inventors for their contribution to future inventions enabled by their idea, smooth licensing markets, and force secrets related to the invention to be disclosed. The basic costs are also fourfold: patent holders can charge above marginal cost, limit diffusion if there are transaction costs in licensing, cause excess spending on races to acquire patents or attempts to "invent around" the patent ex-post, and in some cases distort the direction of invention.

Appropriation of the unique fixed cost paid by the initial inventor is a longstanding justification for patents. If a firm builds a car factory, their rival who wishes to compete will have to incur the same fixed

[62] See https://www.uspto.gov/learning-and-resources/electronic-data-products/patent-claims-research-dataset.

[63] See https://www.patentsview.org/web.

[64] There are alternative ways to protect an idea from competition, however. For example, a trade secret permits a company to attempt to keep proprietary information from the view of the public or competitors, as in the case of the Coca-Cola recipe, and to recover damages from anyone revealing these secrets. The difference is that trade secrets require no disclosure, but also offer no protection from a rival who independently figures out how to build your product.

expense. However, if a firm invents a new engine, its rival can in many cases imitate that idea at much lower cost. Therefore, as John Bates Clark argued, "if an invention became public property the moment that it was made, there would be small profit accruing to any one from the use of it and smaller ones from making it" (Clark, 1915). Even earlier, Jeremy Bentham made an even more pithy argument: patents are necessary so that "what is sown may be reaped," or – in more modern terms – the net social value of an invention will exceed the net private value if the cost of the invention cannot be recouped before competition arrives (Bentham, 1825). On the other hand, inventions that are conceived as a byproduct of normal production, are already protected by first-mover advantage, or would have been invented because of luck or individual curiosity even in the absence of a patent do not require this monopoly right as a mechanism to overcome the inventor's fixed costs. Nordhaus (1969) and Scherer (1972) offer formal analyses of this tradeoff, the latter giving a particularly easy-to-understand geometric argument.[65]

If ex-ante incentives for inventors are important, then so is ex-ante competition. Let a set of inventors have the ability to do research on an invention over time. With some probability depending on their effort, an invention is found, and a patent is granted to the inventor. The patent race model of Loury (1979) models the patent arrival rate as an exponential function of research effort by each given firm. At the start of the game, each firm pays a fixed cost equal to their chosen research capacity, with research effort free thereafter: that is, firms pay for the "size of the lab." In equilibrium, firms overinvest in R&D because they do not account for the fact that their effort lowers the probability rival firms will invent first. Lee and Wilde (1980) modifies Loury's model such that firms only pay for R&D while they are performing research, hence if rivals invent quickly, costs incurred by noninventors are low. This modification changes the relationship between competition and research effort at any given time, but does not flip that conclusion that too much R&D effort will result. Reinganum (1982) and Judd (2003) extend patent races to differential games where firms choose research intensity over time and the value of inventions may vary depending on exogenous features or the research results of other players. All of these models suggest that the ex-ante distortion of patents on researcher effort must be weighed against the ex-post benefit of easier appropriation. By incorporating these ex-ante incentives, they provide a distinct contrast with their dynamic racing externality from "all-pay auction"-type models where the firms that invest the most receive the good with certainty (e.g., Barzel (1968) and Gilbert and Newbery (1982)).

Patents in practice protect not only the direct invention, but also some inventions which build on that initial invention. Green and Scotchmer (1995) notes an intriguing tradeoff: with sequential invention, there is not enough total social value to provide effective incentives on the margin for each inventor. Let one firm have an idea with no social value as a product, but which makes a second invention possible. That initial idea costs c_1 to develop. A second firm has the ability to develop that follow-on invention with social value v at cost c_2. It is optimal to invent both products if $v > c_1 + c_2$. However, the *marginal* social value of each of these two complements is v, hence the sum of marginal values is $2v$. If each firm earns their marginal value, society pays more than the total social value of the combined set of goods. The problem is clear to see when c_1 is close to zero and c_2 is nearly equal to v. If the initial inventor does not receive a patent covering the second invention, she will not earn any revenue and hence nothing will be invented. On the other hand, if the initial inventor has a patent covering the second invention, the second inventor will not bother unless he earns at least c_2. This may not happen if license negotiations occur after that second invention is made, in which case the second inventor will earn only $\alpha \times v$, where α is a bargaining parameter.[66] Optimal sequential invention therefore depends both on early inventors having a patent covering later inventions, and the ability to identify future inventors before they sink R&D costs so that efficient bargaining without holdup can occur.

[65] Note that we here focus on describing the current structure of the patent system, which is intended to be uniform across technologies as opposed to being technology-specific; on the latter, see Burk and Lemley (2003) and Roin (2014b), as well as Acemoglu and Akcigit (2012), which argues that optimal policy will involve state-dependent intellectual property right protection.

[66] Note that patent breadth therefore sets the threat point in bargaining. Indeed, a patent holder may prefer a less broad patent since she cannot commit to not holding up the second inventor, and that commitment problem causes some Pareto-improving inventions to never appear.

Denicolò (2000) combines patent races and sequential innovation, pushing back against the idea that we want to give strong patent rights to initial inventors. Racing for that initial patent will be overprovided in equilibrium, and tilting profits toward those initial inventors disincentivizes effort on follow-on inventions. Gallini (1992) points out that if it is possible to invent substitutes for a patented invention at cost, then long-lived patents can cause more competition in the short-run by inspiring the invention of these competing products. Overall, the theoretical literature on patents clarifies that they not only allow inventors to appropriate the value of an invention, but also affect ex-ante incentives to perform R&D at all, ex-post sequential invention, and the type of invention firms pursue.

The sequential nature of innovation, where IP affects incentives to invent directly today and by changing the nature of competition tomorrow, can lead to some counterintuitive results. In a Neoschumpeterian model, Acemoglu and Akcigit (2012) argues that firms who invent when they have large technological leads should get *stronger* patents, both to incentivize them to keep working and to incentivize laggards to prevent such a large lead from opening. Purely static models with one-shot inventions may miss important incentive effects from endogenous dynamic market structure. The patent tradeoff is therefore much deeper than simply balancing monopoly distortions against the provision of quasirents to incentivize costly research.

An important clarification to the role of patents requires considering the information and commitment power of the planner. As Wright (1983) correctly notes, patents are simply a mechanism for transforming the shadow price of potential inventions into a monetary reward. If the planner knows the social value and expected R&D cost of an invention, then they can simply pay the lowest-cost inventor their expenses and hence induce all inventions with net positive social value. That said, although the planner *can* pay this ex-post, will they? Roin (2014a) argues that the patent system both prevents government from reneging on payouts after an invention appears, and leads to higher rewards to inventors whose invention turns out to be more valuable over time. Patents are therefore useful both because of asymmetric information between inventors and the planner about what should be invented, and because of a commitment problem on the part of government.

Ex-ante, it is likely true that planners do not know which inventions are worth pursuing and how hard they are to pursue. Ex-post, however, it may be possible to reduce the deadweight loss of monopoly pricing by patent holders. Scotchmer (1999), Hopenhayn and Mitchell (2001), Hopenhayn, Llobet, and Mitchell (2006), and Weyl and Tirole (2012) all explicitly consider patent policy as a potential solution to a screening problem in mechanism design. Kremer (1998) proposes a "patent buyout" where patents are sometimes forced into a public auction and the planner sometimes pays a premium on the auction price in exchange for bringing the invention into the public domain. The probabilistic aspects ensure that rival firms with information about the value of the invention have an incentive to bid, and hence that the buyout price does not suffer from a government commitment problem where its incentive to pay falls once the invention exists. Whether ex-post rewards or buyouts are a useful complement to patents given this commitment worry is a debate going back to the nineteenth century: see Shavell and van Ypersele (2001) in partial support, Khan (2015) in opposition, and Galasso (2020) on dynamic reward policies which limit the commitment problem.

The direction of research - that is, which projects are undertaken rather than the aggregate sum of R&D - is also likely to be inefficient even with patents. Nelson (1959b) famously diagnoses a particularly worrying directional distortion: firms will perform too little basic research. He notes that the uses of basic research are broad and hence unlikely to be important to any particular firm, and further than because basic research is hard to embody in a product or technique, it may not be possible to patent and then license to outside firms. That laissez faire markets or even patents will lead to distorted direction is a broader problem than simple underappropriation, however. Dasgupta and Maskin (1987) allows competing firms to choose research agendas which may be correlated with each other. If one firm reduces its correlation with the other, the probability their inventions are wastefully similar is reduced. This externality implies that research portfolios are too similar from a societal perspective. Acemoglu (2012) considers inventions which can be improved in the future, after a patent or first-mover advantage expires. Inventions with little value given consumer preferences today, but which will be built on by inventions tomorrow to satisfy future preferences, are underprovided. Bryan and Lemus (2017) models inventions which vary in their immediate payoff, their difficulty, and the future inventions they make possible. R&D on a given project may be suboptimal either because the inventor gets only a small portion of the full social value their invention creates inclusive of continuation value, or because the invention is difficult and hence rivals race to invent easier substitutes. Because the relative size of these two effects depends on parameters of potential inventions which are never invented in equilibrium, distortions in direction are not wholly fixed with policies like patents or research subsidies. In a similar model, Hopenhayn and Squintani (2021) focuses on dynamic congestion externalities with switching costs across research projects: among other distortions, "hot" technologies attract too many inventors since none weigh the cost of rivals having to redirect research after a successful invention. Akcigit, Hanley, and Serrano-Velarde (forthcoming) studies the choice of basic versus applied research in an endogenous growth model based on Klette and Kortum (2004) calibrated to match French R&D data. They find that firms perform too little basic research, that patent race effects mean applied research is if anything overprovided, and that neutral policies like R&D subsidies will not fix these problems.[67]

Beyond ex-ante and ex-post incentives for inventors, patents also play a role in smoothing the market for technology. It matters not just what is invented, but who is able to commercialize the invention, because agents differ in their capacity for successful commercialization (Arora, Fosfuri, and Gambardella, 2001; Teece, 1986). Beyond capability for commercialization, joint profits are often higher following licensing since the incumbent and the innovator can act as a monopoly rather than compete (Gilbert and Newbery, 1982). However, as Arrow (1962a) explains, firms will be reluctant to disclose their technology to potential partners and licensors if they think the technology will be expropriated. Patents limit this risk by giving the patent holder a legal remedy. Hellmann (2007) notes that patents may also improve technology markets by allowing those with lower costs of searching for technology-user matches to perform this search. That said, the possibility of licensing not only smooths profit-enhancing transfers ex-post, but also can distort incentives on margins which affect bargaining over licenses. For example, Gans and Stern (2000) argues that licensors have an incentive to develop research capability they will not use in equilibrium to improve their outside option if licensing negotiations break down. A firm that can credibly threaten to invent a substitute if they do not get a good deal on a license is in a better bargaining position.

The possibility of using secrecy, whether informal or as part of a trade secret, also affects firm behavior by presenting a choice: receive a patent with strong protection for a limited time, or rely on secrecy with weaker protection but potentially more time.[68] Anton and Yao (2004) explicitly considers this tradeoff in a model of cost-reducing innovations. For minor inventions, patents are sufficient to prevent imitation since the risk of a court forcing the infringer to pay damages is too high to justify the minor benefits of imitation. For big inventions, however, it may be worthwhile not to patent at all. Not getting a patent signals that the inventor is confident they have a substantial technological lead, reducing incentives for rivals to continue competition. As Anton and Yao note, Ford Motor Company actively disclosed partial details of their assembly line process to prove to rivals how far ahead they were technologically, hence to reduce the "neck-in-neck" incentive to compete in Neoschumpeterian models. Scotchmer and Green (1990) notes that secrecy is useful to the dynamic invention process as well. A firm that patents every minor step gives information to rivals which helps rivals more quickly invent high-value inventions which build on those minor steps but which are not fully protected by the initial patent.

A clever argument in Henry and Ponce (2011) shows that the ability to disclose in addition to the ability

[67] In addition to the effect of IP on the direction of research, there is a literature dating back to Hicks on how changes in factor prices affect which inventions arrive. Acemoglu (2002) clarifies a number of longstanding conceptual misunderstandings about this "induced innovation."

[68] See Lemley and Shapiro (2005) on why even patents only provide "probabilistic" protection, as a granted patent with some probability will be revoked by a court if challenged. Schankerman and Schuett (forthcoming) models the patent review process under probabilistic protection, where a government wants to prevent low-quality "obvious" patents from being granted. They suggest more reliance on application fees rather than post-grant fees, as the former deters applications that on the margin are more likely to be turned down on examiner review.

to keep secrets is useful for appropriating rents. Consider an inventor who cannot get a patent. Two other firms, A and B, can either reverse engineer that invention at a cost, or acquire a license. Each firm would like to wait to enter: once their rival knows how to make the invention, the third firm can buy the secret at a price which will be Bertrand-competed down to zero by the two firms who know this information. This desire to wait induces a mixed strategy equilibrium where the initial inventor earns monopoly profits until a random time plus licensing revenue plus triopoly profits thereafter. Henry and Ponce show that these combined profits can exceed the profit earned by a patent holder since the patent length is a fixed, finite time.[69]

In some cases, inventors may wish to *avoid* secrecy, or even to be able to commit to openness. Consider what Allen (1983) calls "collective invention." When a technology is new, incentivizing the development of complements or micro-improvements may be more important for profit than expropriating rents from licensing the undeveloped technology. Consider an iron producer in 19th century England who has an idea on how to improve the efficiency of a blast furnace. Slowing down the collective industry improvement of furnaces by patenting the improvement may lead to lower profit than allowing collective development of blast furnace technology. To put it simply, in early industries, factors which limit the growth of the size of the pie are more harmful to all firms' profits than those which limit the size of the slice any individual inventor can take. Nuvolari and Sumner (2013) describes this process of collective invention, with patents playing a very limited role, in the entertaining case of the invention of porter beer. Bessen and Maskin (2009) investigates collective invention theoretically, arguing that not only are firms willing to do R&D without patent protection in early industries where sequential and complementary invention is important, but that they may be strictly better off if all firms are unable to patent their improvements.

Finally, theorists have long been interested in *who* should enforce intellectual property rules. In particular, should countries enforce patents from abroad, and should poor countries enforce them at all? Intuitively, a country that does little domestic innovation appears to hamstring its ability to catch up to the frontier by denying its firms the right to copy foreign ideas. Prior to the 1883 Paris Convention, patents were largely unenforced outside of their initial jurisdiction (Moser, Bilir, and Talis, 2011). Even after that date, countries frequently made little attempt to enforce foreign patents in technologies where they were trying to catch up to the global frontier. Richter and Streb (2011) shows Germany limited foreign patent rights in the late 19th century and after World War I, times they were attempting to catch up with American toolmakers. In line with this history, Deardorff (1992) argues that, from the perspective of developing countries, the monopoly price harm of patents is balanced against only the *marginal* increase in inventions induced by expanding patent protection to their region: inventors already have some incentive to invent because of monopoly protection in their home country. He therefore suggests that expanding patent protection globally is unlikely to increase welfare. On the other hand, Diwan and Rodrik (1991) cautions that a lack of IP rights in developing countries will distort inventor effort away from technologies that benefit those countries, a problem which may be particularly severe in medicine. Chen and Puttitanun (2005) theoretically models this issue by considering incentives for domestic innovators against the welfare costs of limiting domestic imitators copying ideas from abroad. This work largely focuses on inventive effort as the outcome, rather than the role of IP in the diffusion of new inventions as we will discuss in Section 4.

3.2.3 Intellectual property: Evidence

Let us now discuss what we know empirically about the theoretical tradeoffs described above. Patents are traditionally seen as trading off the benefit of more invention against the cost of higher prices during the life of the patent. The more effective patent laws are at encouraging research investments, the stronger the case for longer patent terms or broader patent grants. A large historical literature has argued that market

[69] This subsection is necessarily a limited review of the broader theoretical literature on IP, which covers topics like the optimal breadth of patents, the role of licensing frictions, complementary assets and appropriation, and more. See Gallini and Scotchmer (2002) and Rockett, Hall, and Rosenberg (2010) for a broader look at the literature.

incentives have played an important role in incentivizing innovation (Khan and Sokoloff, 1993) and that patents played an important role in facilitating the market for technologies (Lamoreaux and Sokoloff, 1999; Lamoreaux, Sokoloff, and Sutthiphisal, 2013). However, when investigating specific policy changes to test for evidence of whether stronger patents induce more R&D, a number of papers have failed to uncover such a relationship (Lerner, 2009; Moser, 2005; Sakakibara and Branstetter, 2001).[70]

A small number of industries stand out as important exceptions to the nonimportance of patents: pharmaceuticals, chemicals, and agriculture. Evidence from surveys (Cohen, Nelson, and Walsh, 2000; Levin et al., 1987; Mansfield, 1986a) as well as data from the pharmaceutical industry (Budish, Roin, and Williams, 2015) are consistent with patents having an economically important effect on R&D in that industry. Moscona (2020) analyzes the introduction of patent protection for plant biotechnology in the US in the 1980s, which affected crops differentially depending on their reproductive structure. He documents compelling evidence that the introduction of patent rights in this sector increased the development of novel plant varieties, and that US counties more exposed to the policy change by nature of their crop composition saw agricultural land values and profits increase. Hence, while skeptics of the patent system such as Boldrin and Levine (2013) rightly point out the general lack of evidence on patents encouraging research investments, the types of analyses presented in Budish, Roin, and Williams (2015) and Moscona (2020) suggest that more work is needed on this key question.

The prediction that stronger patent protection induces additional research investments emerges unambiguously from a class of theoretical models that treat innovations as isolated discoveries. However, in practice – as discussed in Section 3.2.2 – innovation is often "cumulative," in the sense that any given discovery is also an input into later follow-on discoveries. In such cases, optimal patent policy depends on how patents affect sequential innovation.

There is some evidence that intellectual property rights including patents (Galasso and Schankerman, 2015) and copyright-type agreements (Williams, 2013) can hinder follow-on innovation, although a case study of gene patents found no evidence for such an effect (Sampat and Williams, 2019). These results are collectively in line with the theory in Green and Scotchmer (1995): well-defined property rights for an invention and the ability to identify follow-on inventors before they sink R&D costs, conditions likely to hold in the case of gene patents, can lead to efficient sequential innovation. Patents are more likely to be harmful when the precise follow-on inventions they cover are unclear, or when follow-on inventors are worried about holdup from "submarine patents." The question of how patents affect follow-on innovation is particularly important as an area for future work given that the assumption that patents hinder follow-on innovation has informed a recent set of important US Supreme Court decisions restricting the set of discoveries eligible for patents.[71]

Finding appropriate data to analyze the effects of patent policy has been challenging. As discussed in Section 1.3.1, researchers in economics and related fields frequently use patents as a measure of innovation. However, changes in patent law often affect the costs and benefits of filing for patent protection, and hence the propensity to patent a given invention. Changes in patenting and changes in real innovation are therefore conflated. For example, Abrams (2009) investigates the 1995 Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS), which induced patent term extensions that varied across technologies as a function of historical USPTO review times. He tests whether technology classes with longer patent term extensions saw increases in patent counts, but such a test is hard to interpret given that the policy change altered incentives to file for patents on existing research investments.

There have been at least two clever solutions proposed to address this concern. Lerner (2009) analyzes changes in patenting behavior of inventors both in a country experiencing a patent law change *and in foreign*

[70]See also Budish, Roin, and Williams (2016), which notes that many of the empirical approaches used in this literature may be under-powered to detect an effect.

[71]We here refer to Bilski v. Kappos on patents on abstract ideas (investment strategies), Mayo v. Prometheus on patents on laws of nature (diagnostic methods), AMP v. Myriad on patents on natural phenomena (human DNA), and Alice Corp v. CLS Bank on patents on abstract ideas (software).

markets where no such measurement challenge arises; his findings suggest little response of innovation to strengthened patent laws. Galasso and Schankerman (2015) analyzes how follow-on innovation as measured by patent citations change when a focal patent is invalidated, the idea being that even invalidated patents qualify as prior art. Citations to invalidated patents rise substantially following invalidation, with substantial heterogeneity across industries, suggesting patent-related frictions which harm sequential innovation. That said, in our view, most progress in this space has come from developing linkages between patent data and real-world outcomes, such as Moser, Ohmstedt, and Rhode (2018) on the yields of plants covered by plant patents, Murray and Stern (2007) on citations to scientific papers disclosing the same knowledge as a given set of patents, or Sampat and Williams (2019) – building on the work of Jensen and Murray (2005) – linking patents with the genes they claim as intellectual property. More work developing such linkages, such as through the IPRoduct project,[72] would be very useful.

Finding appropriate sources of empirical variation to analyze these questions has also been challenging. On paper, the patent system actively strives to provide uniform protection – a 20-year term that is constant across technologies and across (most) countries. Historically, there did exist some cross-country variation in patent protection, and indeed that type of variation has proven useful in some contexts such as testing how country-level patent protection affects country-specific drug launch decisions (Cockburn, Lanjouw, and Schankerman, 2016; Duggan, Garthwaite, and Goyal, 2016; Kyle and Qian, 2014) and in testing how country-specific patent laws affect the direction of technological change within a country (Moser, 2005). However, because inventions tend to be developed for a global market, a priori it is challenging to use patent law changes in one "small" country to investigate how patent laws affect innovation, given that such a policy change likely has little effect on global incentives for innovation. Moscona (2020) is one example of how single country case studies can be leveraged in the case of large market countries.

The literature on how patents affect follow-on innovation has recently seen the development of several new empirical strategies leveraging administrative data. Galasso and Schankerman (2015) studies patent invalidations and follow-on innovation, as discussed above, by leveraging the quasi-random assignment of judges to cases at the US Court of Appeals for the Federal Circuit. Similarly, Sampat and Williams (2019) – building on the work of Cockburn, Kortum, and Stern (2003) and Lemley and Sampat (2008) – leverages the quasi-random assignment of patent applications to patent examiners at the USPTO to investigate how (instrumented) patent grants affect follow-on innovation. This patent examiners empirical approach has been subsequently refined and improved by work including Feng and Jaravel (2020) and Righi and Simcoe (2019).

Let us briefly mention a number of other patent policy issues that we view as ripe for future work. Only a small number of papers, such as Sakakibara and Branstetter (2001) and Lerner (1994), have empirically investigated patent scope even though scope is clearly a central feature of the patent system. A number of papers in the law and economics literatures have investigated particular aspects of the disclosure function of the patent system (Furman, Nagler, and Watzinger, forthcoming; Ouellette, 2012; Roin, 2005), but relatively little is known from an empirical perspective about the overall costs and benefits of disclosure. One exception is a number of recent papers which have investigated the effects of the 1999 American Inventors Protection Act (AIPA), which required US patent applications to be published regardless of whether or not they were granted a patent, normally 18 months after the application date; Lück, Balsmeier, Seliger, and Fleming (2020) suggests AIPA led to less duplication of invention, and Hegde and Luo (2018) estimates that AIPA induced a reduction in time to licensing. Policy interest in alternatives to patents – such as prizes, contests, patent buyouts, and laissez faire alternatives such as secrecy and first mover advantage – have been the focus of some excellent work (Anton and Yao, 1994; Gallini and Scotchmer, 2002; Hopenhayn, Llobet, and Mitchell, 2006; Khan, 2015; Kremer, 1998; Moser, 2012; Teece, 1986; Wright, 1983), but nonetheless seem ripe for further development. Finally, a number of papers have investigated the role of excludability and open access – separate from patents – and provided some facts laying the groundwork for future work as well (Bryan and Ozcan, forthcoming; Cohen and Walsh, 2007; Furman and Stern, 2011; Murray, Aghion, Dewatripont,

[72] See https://iproduct.io/app/#/public/page/home.

Kolev, and Stern, 2016; Murray and Stern, 2007; Walsh, Cho, and Cohen, 2005; Walsh, Cohen, and Cho, 2007; Williams, 2013).

3.3 Competition policy

The rate and direction of innovation both determines and is endogenous to market structure. Indeed, Aghion and Tirole (1994) calls this relationship the second-most studied question in all of industrial organization, after the link between market structure and profits. Antitrust, and competition policy more broadly, is therefore an important policy lever for properly incentivizing innovation. That said, a reason the link between market structure and innovation is so heavily studied, and so contentious after all of this study, is that the relationship is both theoretically complex and empirically challenging to measure.

The earliest economic argument for why competition policy should care about innovation came from Joseph Schumpeter. He famously described dynamic competition, whereby large corporations bring new products to market through costly effort, as the "most powerful engine for progress" under which "perfect competition is inferior" since all profits which might fund research are competed away (Schumpeter, 1942). Implicit in this argument is that recovering the fixed costs of innovation requires either sufficient scale or profitability per unit. It is hard to overstate the import of this argument. Schumpeter is arguing that Marshallian welfare analysis, where the "best" market structure is perfect competition that drives markups to zero, is exactly backward on how we should regulate innovative industries. If price falls to marginal cost, there are no rents to pay for innovation, and hence no incentive to innovate.

Theoretically, however, things are not so simple, even restricting to models where the incentive to innovate depends only on the expected return to innovators. Arrow (1962a) considers the inventor of a cost-reducing invention who can make a take-it-or-leave-it offer to producing firms. For a small cost reduction of Δ, the inventor will charge all firms in the competitive market precisely Δ per unit produced, earning Δ times the pre-invention competitive quantity sold. If selling to a monopolist, however, the profit and hence incentive to invent will be less: to reduce costs, the monopolist is willing to pay no more than Δ times the *monopoly* quantity sold. Competitive industries therefore see more innovation, not less. More generally, since monopolists care about how an innovation will raise their pre-innovation profits, and competitive market innovators care about the overall value of that new innovation, the "replacement effect" means monopolists will innovate less.[73]

It gets worse. In their famous "inverted-U" model, Aghion, Bloom, Blundell, Griffith, and Howitt (2005) explicitly adds dynamics, finding yet another theoretical link between structure and innovation.[74] When firms in a given market have similar technology, then small inventions are not very profitable: they neither make the inventing firm's product so superior that a high price can be charged, nor do they cause the firm to have such a large technological lead that rivals give up trying to catch up with their own inventions. When one firm has a large technological lead, likewise there is little incentive to innovate, as the firm's existing product is likely quite profitable already and laggards realize they have little chance of matching the leader's technology. It is the intermediate range, where Schumpeterian profits are high due to innovators having some market power, but also where preinnovation rents are low, encouraging these firms to try to "pull away" from competitors, that innovation will be maximized. This intermediate range of competitive parity is, in dynamic equilibrium, equivalent to a market structure with intermediate industry profits.

[73] And even here we must be careful. Greenstein and Ramey (1998) notes that for product innovations, an existing monopoly can use a new differentiated product to sort customers by willingness to pay, earning more marginal profit than an entrant innovator. That is, it is not the case that new products necessarily replace existing sales - they can boost them. The same is, of course, true of products that are complements to the incumbent's existing products. Similarly, Gilbert and Newbery (1982) argues that when the new product competes with the old, the incumbent weighs monopoly versus duopoly profits when trying to win the patent race, and the entrant weighs duopoly versus zero profits. The relative ranking of these two differences is ambiguous in general. See also Reinganum (1985), Harris and Vickers (1985) and Doraszelski (2003) for subtleties of incumbent-entrant competition.

[74] The inverted-U hypothesis also appears in a slightly different model in Lee (2005), and was first noticed empirically in Scherer (1967). See also Aghion and Griffith (2008) on pre- and post-innovation rents and the incentive to innovate.

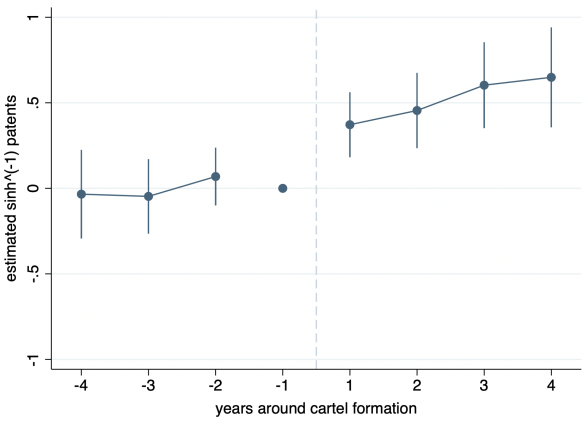

(a) Collusion and innovation (Kang, 2020)

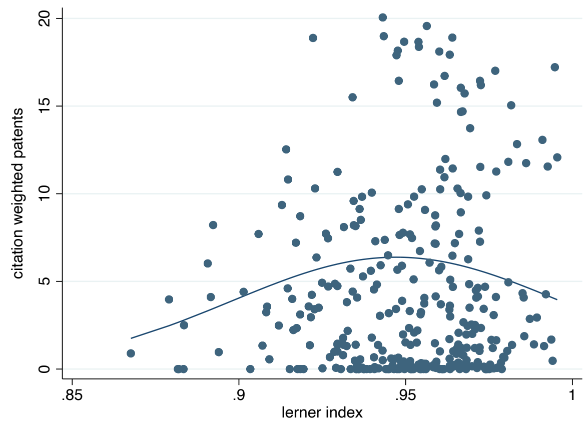

(b) Competition and innovation (Aghion, Bloom, Blundell, Griffith, and Howitt, 2005)

Figure 10: The relationships among collusion, competition, and innovation

Notes: Panel (a): Following a change in market structure related to the formation of a cartel later uncovered by US antitrust authorities, citation-weighted patenting activity increases 57.7%. The plotted estimates are event-time coefficient estimates of the inverse hyperbolic sine transformation of citation-weighted patents by firms before and after a cartel is formed. Panel (b): This figure plots the Lerner index as a measure of competition on the x-axis against citation-weighted patents on the y-axis. Each point represents an industry-year. Only data points that lie between the tenth and ninetieth deciles of the citation-weighted patents distribution are included in this figure. The exponential quadratic curve that is overlaid was created by a Poisson regression of citation-weighted patents on competition and competition squared, including year and industry fixed-effects.

It seems we are at an impasse. Innovation may be greater in monopolized industries with high post-innovation rents (Schumpeter), in competitive industries where the total quantity of products to be improved is higher (Arrow), or in intermediate settings where the dynamic gap between pre- and post-innovation profits is highest (Aghion et al). Figure 10 displays an example of this contrasting evidence: the left panel plots increased patenting following a reduction in competition after cartel formations, while the right panel shows an inverted-U pattern in a measure of competition against the number of citation-weighted patents per firm.

Each of these theories implicitly build in assumptions about imitability, appropriation, and the ability to sell inventions onward to other firms. The fundamental tension of innovation antitrust is whether market power in innovative industries is a sign of quasirents recovering past R&D spending ("competition for the market") or of inefficient barriers to entry permitting supernormal profits for technological leaders. Indeed, even these barriers to entry may be the result of past innovations: firms endogenously choose to incur some fixed costs which change the scale at which competitors can enter (Sutton, 1998). As these "endogenous sunk costs" are often themselves innovation, the activity a firm with market power pursues in order to prevent potential future competition affects the market structure-innovation relationship. Theory in these broad strokes offers few clear tests which can be applied to existing industries or proposed mergers.

Empirically, competition policy in innovative industries is even more complex, for three reasons. First, as we have noted many times in this chapter, measuring innovation is challenging: ought we use patents, or citation-weighted patents, or direct measures of R&D inputs, or some alternative? Should we be weighting innovation by firm size, when firm size is itself endogenous to market structure?[75] Second, the appropriate measure for competition is not clear: Herfindahl indices, a Lerner index, or a measure of business dynamism, all of which can be measured among technologically-similar firms or among downstream competitors? Third, since innovation has a first-order impact on market structure, pure correlational analyses are likely to mislead.

With these caveats in mind, early empirical literature on market structure and R&D generally found no

[75] See Sutton (1998) for a basic argument about sunk costs and free entry.

relationship between the two. Gilbert (2006) provides a thorough review of these early studies, noting that most treated Schumpeter's hypothesis as an argument that bigger firms do more R&D, rather than the more precise statement that more concentrated markets do so.[76] An important caveat is that raw cross-sectional analysis is fundamentally limited by the variance of technological opportunity across industries. For instance, Geroski (1990) notes that including industry fixed effects in a panel of British data completely overturns an uncontrolled result that more concentrated industries are more innovative. Sutton (1998) argues theoretically and empirically that an analog of minimum efficient scale which varies across industries - the homogeneity of submarkets to which a new product can be targeted - endogenously affects the number of firms who will find it worth doing R&D, and hence endogenously affects market structure.

That said, even recent studies with more sophisticated reduced-form empirical methodologies do not find an unambiguous monotonic relationship between market structure and innovation. Kang (2020) uses the breakup date of all known US collusive cartels to investigate whether more concentrated industries innovate more or less. He finds that while price-fixing is in effect, cartel members file 48% more patents and 33% more top quality patents, do 18% more R&D, and patent in 30% more technology classes, relative to otherwise-similar matched firms. Using the differential shock of Chinese export-growth across US industries, Autor, Dorn, Hanson, Pisano, and Shu (2020) finds that firms facing more competition do less R&D and produce fewer patents, controlling for industry-level trends that predate modern Chinese export growth. On the other hand, Lampe and Moser (2016) finds that collusion on the R&D side in the form of shared patent pools decreases overall R&D intensity. These results are not necessarily in conflict: price collusion or insulation from foreign competition raises industry returns, while input collusion lowers an individual firm's return to R&D.

Beyond affecting the rate of innovation through post-entry profits, competition may also push X-inefficient firms to pursue frontier technology (Leibenstein, 1966). For instance, competitive industries cause inefficient firms to fail (Syverson, 2004), and drive firms to adopt better management practices (Bloom and Van Reenen, 2007). There is also a small literature arguing that moderate levels of competition inspire higher-variance invention. Gross (2020) uses variation in the level of competition in logo design contests to show that while low post-entry expected profits cause "R&D" effort to fall, a high probability of winning the logo contest causes firms to design very conventional logos. Therefore, intermediate competition among innovators produces the most creativity. This is unsurprising: in winner-take-all contests like many R&D races, high variance research is the only type that can win (see also Cabral (2003)).

Although the general link between competition and innovation is empirically ambiguous, theory and structural models are sometimes useful for examining specific policies or the quantitative magnitude of offsetting tradeoffs.[77] For instance, Segal and Whinston (2007) asks whether innovation is dulled if we let incumbents sign exclusive contracts with buyers. On the one hand, being an incumbent is now more valuable, while on the other hand, entrants have less incentive to do the R&D necessary to enter this market. However, a fairly straightforward argument suggests that aggregate innovation rises when exclusive buyer contracts are limited, as does welfare, by limiting the fraction of customers who are forced by contract into old technology even though better alternatives have been invented. That is, the complexity of the relationship between R&D and market structure does not mean we have to punt on all antitrust questions of interest.

More recently, a series of papers have estimated dynamic structural models of R&D, where differential costs and incentives to innovate are identified from forward-looking behavior, allowing the causal relationship between policy-driven changes in market structure to be recovered from data. Goettler and Gordon (2011)

[76] Of course, large firms may have higher returns to R&D not only because they face less competitive product products, but also because of cheaper internal cost of capital, economics of scale and scope in research production, more diversified risk, or better ability to spread fixed costs across more units of final product. See Cohen (2010) for a detailed review of this literature showing, broadly, that R&D rises proportionally with firm size, and that its productivity falls. Cohen and Klepper (1996) offers a clever argument and supportive evidence from business line data that this pattern is due to cost-spreading: large firms spread R&D over more final sales, hence do R&D on the margin that is less productive.

[77] See Shapiro (2012) and Baker (2007) on where most models agree as to appropriate competition policy, with application to a series of recent proposed mergers in the pharmaceutical industry in the former.

models competition between AMD and Intel dynamically, suggesting that Intel would have innovated more than the duopoly had it been a monopolist due to the higher prices it would have been able to charge. On the other hand, Igami (2017) finds that 57% of the gap in innovation between incumbent and entrant hard disk drive manufacturers is due to Arrovian replacement effects dulling the incentives of market leaders. Igami and Uetake (2020) solves a dynamic structural model of competition and innovation in the same industry. Dynamic incentive to preempt competition causes innovation incentives to increase steeply as we move from one to three or more competitors. Yang (2020) studies the incentives to innovate in a vertical relationship, such as System-on-a-Chip vendors selling technology to cell phone manufacturers. Empirically, the benefits of a hypothetical merger between Qualcomm and a smartphone manufacturer, who can then coordinate R&D incentives and reduce double marginalization, exceed the harms from foreclosing entry or increasing costs for other smartphone makers.

Beyond generic merger policy for innovators, a particular competition issue which has attracted substantial policy attention in recent years is the desirability of acquisitions by technological leaders in innovative industries. In an influential paper, Cunningham, Ederer, and Ma (2021) argues that "killer acquisitions" may occur, where leaders acquire inventors of competing products and then mothball them in order to blunt competition. Using data on pharmaceutical acquisitions, measuring substitutability by the overlap in method of action and target of the acquired therapy compared to the acquirer's existing portfolio, they find that competing drugs are almost 30% less likely to be developed than non-competing drugs which are acquired. In aggregate, they estimate at least 6% of pharma acquisitions are made to limit competition.

These observations have led to a large and growing set of studies on precisely when and where these acquisitions are worrisome, particularly given the frequency with which small, fast-growing competitors are acquired in high-tech industries (Cruise by General Motors, Waze by Google, Instagram by Facebook, Slack by Salesforce, and so on). These tech acquisitions occur in a very different appropriability environment from innovation in the pharma market. Potential post-merger inefficiencies must be balanced against the fact that the possibility of acquisition by a firm with more suitable complementary assets incentivizes innovation by small firms.[78] Bryan and Hovenkamp (2020) models fringe firms who can sell to technological leaders or laggards, or license to both. In laissez faire, new technology diffuses too narrowly, innovating firms target technologies most useful to incumbent leaders rather than to welfare more broadly, and the ensuing technological gap eventually retards the overall rate of innovation. Simple mandatory licensing rules solve this problem while leaving welfare-enhancing acquisitions unaffected.[79] Kamepalli, Rajan, and Zingales (2020) argues that when network effects and switching costs are prevalent, consumers will not adopt products in the "kill zone" if they worry that large incumbents will acquire and kill the product if it becomes modestly successful. They show that venture capital investment in areas related to firms acquired by Google and Facebook fall substantially after the acquisition, but investment in firms acquired by less dominant players goes up. Cabral (2018) notes that mergers increase the relative return to startups of incremental Arrovian invention relative to incumbent-replacing radical Schumpeterian invention, implying that startup acquisition policy matters for the direction of research.[80]

Two final subtleties for competition policy involve the difference between *technology* competition and *product market* competition, and the role of firm architecture. As Bloom, Schankerman, and Van Reenen (2013) cleverly exploits in its study of spillovers, firms may do R&D in similar areas even though they produce products sold in very different markets: a novel method of aerosolizing a liquid has applications in medicine and in mosquito spray. How licensing and acquisition policy affects innovation in these settings is essentially

[78]Indeed, it may overincentivize them - an entrant may create a product solely to threaten lowered post-entry industry profits and therefore induce acquisition by an existing monopolist (Rasmusen, 1988).

[79]This caveat is important. For instance, Phillips and Zhdanov (2013) finds that increases in industry-level M&A activity increase R&D of small firms, their probability of being acquired, and the share of R&D they perform. That is, many mergers simply reflect division of labor into technology firms and sales firms. As to why *all* startups do not sell to incumbents in order to reduce competitive price pressure, see Gans, Hsu, and Stern (2002) on how dynamic capabilities for innovation make this less jointly profitable.

[80]See also Callander and Matouschek (forthcoming) and Letina, Schmutzler, and Seibel (2021), both of which examine the theoretical conditions under which killer acquisitions will occur and harm welfare.

unstudied. Two exceptions are Marshall and Parra (2019) and Watzinger, Fackler, Nagler, and Schnitzer (2020). The former considers how competition affects innovation when firms are competing both to invent and in the product market. The latter shows empirically that government-mandated compulsory licensing at Bell Labs in the 1950s led to a substantial increase in forward citations to affected Bell patents.

An even more subtle area of inquiry involves adapting the managerial strategy literature on "architectural disruption" to industrial organization. In these models, firm structure, including relational contracts, are sticky, making it challenging for incumbents to adopt technology which does not match their existing capabilities (see, especially, Henderson (1993) and Helper and Henderson (2014)).[81] This stickiness naturally affects the nature of competition for inventions which strong incumbents find difficult to adopt. Incumbents cannot simply set up subdivisions who replicate entrants because of "diseconomies of scale" where firm-wide characteristics, such as a reputation for reliability, cannot be severed between the firm as a whole and the "innovative" subdivision (Bresnahan, Greenstein, and Henderson, 2010). Better understanding of differences in the ability to incorporate innovations into incumbent firms is greatly needed.

To sum up, the question of which market structure maximizes innovation is poorly defined without further assumptions about what type of conduct will be allowed by firms. Indeed, as Gilbert and Newbery (1982) makes clear, if interaction among firms is unrestricted, there is necessarily a negative relationship between competition and profitability: firms that can collude in an unrestricted way will maximize joint surplus. Without knowing what will happen dynamically to the rents of incumbents and entrants, the extent to which the latter can be acquired by the former, the heterogeneity of technological opportunity, and the importance of complementary assets including scale and scope to appropriability, there is little either theoretically or empirically that holds up across industries. Nonetheless, for a given policy environment, industries where the gap between pre- and post-innovation returns for potential innovators is small will see less innovation than those where it is large. Therefore, industries with technologically-dominant firms who sell at monopoly markups will see little innovation, but so will industries where innovators can earn little return due to fierce post-entry competition or limits on acquisition of successful innovators.

Despite these challenges, innovation is so fundamental to economic growth that research showing large benefits on innovation from better competition policy is in high demand.[82] The burgeoning literature on acquisitions whose sole intent is to reduce competition is important, but potential dynamic suboptimality in laissez faire market structure goes well beyond those acquisitions. Convincing empirical evidence of specific changes in competition policy with first-order welfare benefits remains an important open problem for future research. In addition, just as we saw that a justification for patents rather than alternatives like prizes lies in aligning rewards with market sales, and hence solving an information problem for the policymaker, permitting market power for innovators may be a second-best policy when the planner does not know which inventions are ex-ante valuable.[83]

3.4 Labor market policies

Labor supply is a key input into the innovation process. In addition to policies like taxation, policies directly targeting labor supply can affect whether and how an individual will innovate. The supply side of labor, including general human capital development, is of course driven by many factors.[84] A systematic review of all of these factors is beyond the scope of this paper. We therefore focus attention in this subsection on one set of labor market policies – immigration policies.

Why immigration? As we discussed in Section 1, knowledge spillovers are central in innovation policy,

[81] The need to overcome relational contract stickiness as a limit to innovation may help explain why industries insulated from competition appear stagnant. See, e.g., Syverson (2011) on modern treatments of this "X-inefficiency."

[82] See, for example, Tirole (2020), Zingales, Rolnik, and Lancieri (2019), and Crémer, de Montjoye, and Schweitzer (2019) for recent reports on how to modify competition policy for high innovation, digital industries.

[83] de Rassenfosse and Zhou (2021) shows that patents do in fact cause higher markups, especially in competitive product markets and for "important" inventions.

[84] On inequality influencing the supply of innovators, see our discussion in Section 5.

and a variety of evidence suggests that knowledge spillovers have a local component. This simple insight has a variety of important implications, one of which is that it matters not just *who* is innovating but also *where* that person is innovating, since the productivity of a given individual as well as the spillovers of their innovative activity will depend on their local environment. Hence, immigration policies which affect both the propensity of individuals to innovate and the spillovers generated by their innovation are a key policy lever. At the end of this section, we briefly discuss a few other areas of related research. But let us say explicitly that this choice of focus on immigration policy largely reflects the paucity of available research on many other important dimensions of the relationship between human capital policies and innovation.

Historically, the US has had a relatively open immigration policy – a fact which is widely believed to have supported the country being a magnet for both educating and employing foreign-born individuals in science and engineering fields. Looking over the last few decades, a variety of metrics suggest that the share of US-based foreign-born individuals engaged in science and engineering has been increasing.[85] At the undergraduate level, the number of temporary visa holders earning bachelor's degrees in a science or engineering field has more than doubled in recent years, with roughly 15,000 of these degrees being awarded in 2000 compared to 32,000 degrees in 2015. A similar trend appears at the graduate level, with the number of foreign-born students enrolled in a science or engineering graduate-degree program increasing from 128,000 in 2000 to just over 240,000 in 2015. Figure 11 gives a rough sense – at the field level – of what percent of bachelor's degrees and doctoral degrees were earned by temporary visa holders between 2000 and 2015, making clear that the education and training of these individuals is substantial across essentially all fields of study but particularly in various sub-fields of engineering as well as economics.

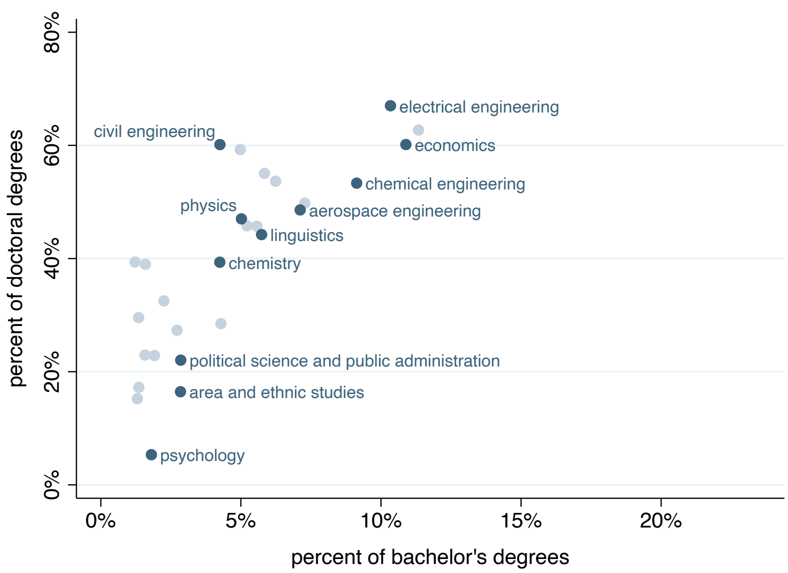

Figure 11: Percent of degrees earned by temporary visa holders, 2000 - 2015

<u>Notes</u>: The x-axis plots the percentage of total bachelor's degrees by sub-field that are earned by students on temporary visas between 2000 and 2015. The y-axis plots the same for doctoral degrees. Source: Appendix Tables 2-22 and 2-32 of National Science Foundation (2018).

From a workforce perspective, the importance of US-based foreign-born workers in science and engineering occupations has likewise been growing in recent decades. Among US-based foreign-born workers who hold a bachelor's degree as their highest degree, only 11.4 percent worked in a science and engineering occupation in 1993; by 2003 this figure had reached 17.5 percent, and it rose further to 21.2 percent in 2015. Similarly, among US-based foreign-born workers with a doctorate, the percent of people in a science and engineering

[85] The statistics quoted in this and the following paragraph are drawn from National Science Foundation (2018).

occupation rose from 26.8 in 1993 to 43.8 in 2015.

Looking beyond these statistics on education and employment, immigrants have long been perceived as generating a disproportionate share of US inventions. Because traditional datasets measuring innovative activity – such as USPTO administrative patent records – do not record an individual's immigration status, this perception has largely been based on anecdotes or highly selected populations[86] rather than on large-scale representative samples. A series of papers by Bill Kerr and co-authors – such as Kerr (2008a) and Kerr and Lincoln (2010) – use ethnic name matching procedures to proxy for immigrant status, but as stressed by Kerr (2008b) this methodology cannot distinguish foreign-born individuals from US natives with ethnic names, and also cannot identify immigrants from Western Europe. Hunt and Gauthier-Loiselle (2010) and Hunt (2011) use survey datasets measuring immigration status, but these of course are limited to relatively small samples.

In a key measurement advance, Bernstein, Diamond, McQuade, and Pousada (2019) leverages a new linkage between the Infutor data and USPTO patent records to sharpen our understanding of immigrants' contribution to US innovation.[87] The Infutor data provides – among many other variables – year of birth as well as the first five digits of Social Security numbers (SSNs) of most adults living in the US over the last 30 years. The authors' methodological advance is to use these two variables to infer immigrant status, given that the first five digits of an SSN pin down the year the SSN was assigned and the fact that in recent decades practically all US natives are assigned a SSN during their youth (often at birth or when they obtained their first job) whereas individuals receiving an SSN in their twenties or later are highly likely to be immigrants. They validate this novel measurement approach using data from the Census and American Community Survey. Linking inferred immigrant status to USPTO records based on individual names and addresses then allows them to tabulate various measures of immigrants' contribution to US innovation. Figure 12, which replicates Figure 2 of their paper, summarizes their key results. In their data, immigrants account for around 10 percent of the population, but comprise 16 percent of inventors and account for around 22-24 percent of innovation as measured either by patents or by a variety of measures that weight patents by their impact as measured by citations or market value. These quality-weighted patent measures reject the view that immigrants are producing more patents of lower quality than US natives.

From a policy perspective, current debates tend to focus on student, J-1, H-1B, and L-1 visas.[88] H-1B visas are granted to people working in a specialty occupation, and they require a higher education degree or its equivalent. J-1 visas are designed for work- and study-based exchange visitor programs. L-1 visas are granted to intracompany transferees in roles requiring specialized knowledge. Figure 13 plots the number of H-1B, J-1, and L-1 visas granted each year in the US from 1987 to 2019. Over this period, the largest increase was in J-1 visas. H-1B visas have increased at a similar but slightly lower rate; notably, over this period a cap of 65,000 H-1B visas was in place, implying that most of the observed growth (to nearly 200,000 H-1B visas granted in 2019) was driven in H-1Bs granted to exempt organizations such as institutions of higher education, non-profit research organizations, and government organizations.

As one of the primary mechanisms through which immigrants in science and engineering fields can enter the US to work, the H-1B visa program has been a focus of both academic work and policy interest. Some

[86]For example, Stephan and Levin (2001) documents that immigrants are disproportionately represented across groups of individuals making exceptional scientific contributions, such as election to the National Academy of Sciences.

[87]Akcigit, Grigsby, and Nicholas (2017a) offers a historical perspective on this question. Although our focus in this chapter is not on entrepreneurship, it is also worth mentioning the work of Azoulay, Jones, Kim, and Miranda (forthcoming) which classifies entrepreneurs as US born or immigrant based on country of birth as recorded in the Census Numident data, in order to document how often immigrants start companies, how many jobs these firms create, and how firms founded by native-born individuals compare.

[88]It is worth noting that although somewhat less directly policy relevant, recent years have seen the emergence of a variety of more historical evidence suggesting that reductions in US immigration damaged US innovation (Burchardi, Chaney, Hassan, Tarquinio, and Terry, 2020; Doran and Yoon, 2020; Moser and San, 2019; Moser, Voena, and Waldinger, 2014). Also notable are Gaulé and Piacentini (2013) which documents that Chinese chemistry PhD students produce much more scientific output during their thesis than other students, and Agarwal, Ganguli, Gaulé, and Smith (2021) which estimates the impact of US immigration barriers on global knowledge production.

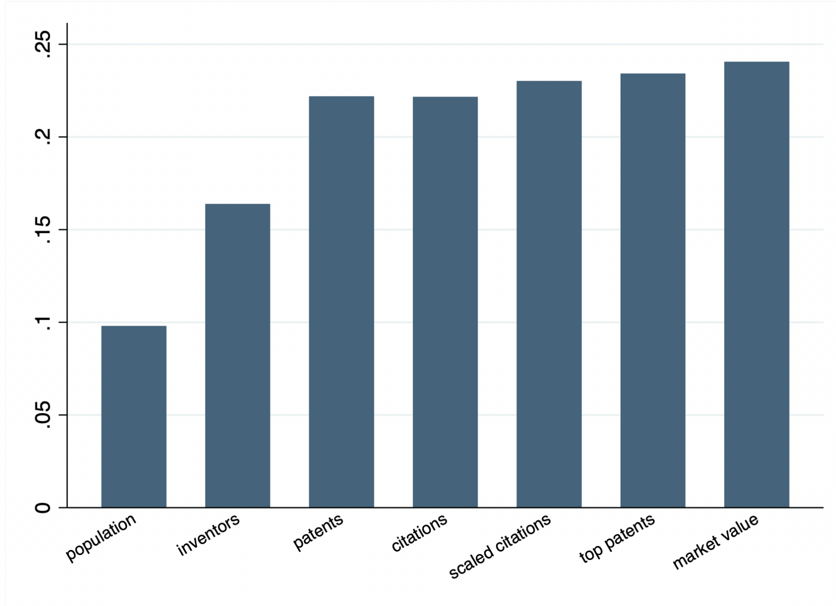

Figure 12: Immigrants' contribution to US innovation

Notes: Categories are: (a) share in the overall population from 1990-2015 according to the American Community Survey (ACS); (b) share of overall number of inventors between 1976-2012, where inventor is defined as an individual who patents at least once; (c) share of overall number of patents at the USPTO from 1976-2012; (d) share of overall number of forward citations of these patents, calculated over a three year horizon to avoid truncation issues; (e) citations normalized by the average number of citations in a given technology class year (the year in which all patents were applied); (f) share of top patents, where a top patent is defined as a patent that is in the top 10% of citations in a given technology class and year; (g) share of patent value, calculated based on stock market reaction to patent approval using the Kogan-Papanikolaou-Seru-Stoffman (KPSS) measure developed in Kogan, Papanikolaou, Seru, and Stoffman (2017), which is available for publicly traded firms only. Source: Bernstein, Diamond, McQuade, and Pousada (2019).

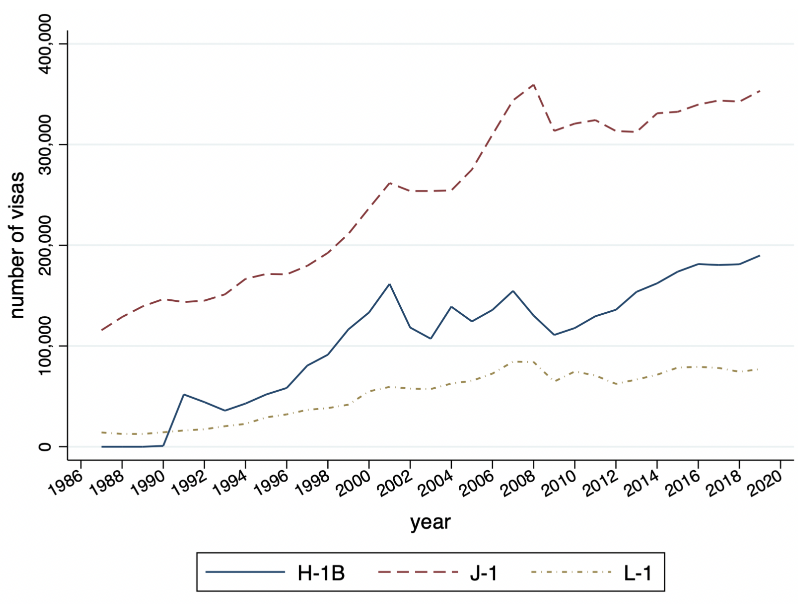

Figure 13: H-1B, J-1, and L-1 visas granted in the US, 1987-2019

Notes: This graph plots the number of H-1B, J-1, and L-1 visas granted in the US each year from 1987 to 2019. Source: US Department of State (2019).

research on H-1B visas has focused on the classic labor economics question of how these immigrants impact native workers.[89] For example, Peri, Shih, and Sparber (2015) uses a shift-share style empirical approach leveraging the share of STEM workers who are foreign in each US city as of 1980, interacted with the over-time change in the number of H-1B visas allocated to different foreign nationality groups in proportion to the city-level presence in 1980. They argue based on this approach that H-1B workers raise native wages and total factor productivity in US cities. In the same *Journal of Labor Economics* issue, Kerr, Kerr, and Lincoln (2015) asks how skilled immigrants like H-1B visa workers impact the employment structure of US firms, arguing that overall employment of skilled workers increases with skilled immigrant employment by the firm. Thus, although the public discourse on H-1B visas often assumes the program harms native workers (e.g. Matloff (2003) and Hira (2010)), the available evidence suggests otherwise.

More recently, a number of papers have investigated how the H-1B visa program impacts innovation by US-based firms. Kerr and Lincoln (2010) uses a shift-share design similar to that in Peri, Shih, and Sparber (2015), and argues invention increases with higher H-1B visa admissions – primarily through the direct contributions of immigrants. On the other hand, Doran, Gelber, and Isen (2014) leverages the sample of firm-worker pairs subject to the H-1B visa lotteries (when the 65,000 cap is hit) and argue that among firms in that sample who win the lottery there is no effect on patenting nor on claiming of the research tax credit. This difference in results could be driven by the difference in samples, or based on a more conceptual difference: the Doran, Gelber, and Isen (2014) study estimates the effect of an additional H-1B visa to one firm on outcomes at that firm, holding constant H-1Bs given to other firms; in practice, crowded-out workers may find employment elsewhere, and innovation could increase at other firms. More recently, Khanna and Lee (2020) links H-1B records to retail scanner data and argues H-1B visas are associated with higher product reallocation and revenue growth at the firm level. Glennon (2020) documents novel evidence that restrictions on H-1B visas cause multinational firms to substitute ("offshore") employment and research investments to foreign affiliates (particularly in Canada, India, and China). Particularly given recent improvements in access to administrative data on the H-1B visa program, this is an area that seems ripe for future work.

Returning to the more general theme of labor market policies, one literature worth highlighting is that of non-compete policies. Marx and Fleming (2012) defines non-competes as employment contracts in which an employee pledges not to work for a competitive firm for a period of time after resigning or being terminated. The classic framing question in this literature is why Silicon Valley is in California rather than Massachusetts (e.g. Gilson (1999)). While impossible to derive a precise answer to that question, many have argued that one candidate explanation is the fact that California law prohibits post-employment non-compete covenants whereas Massachusetts has historically enforced them. Non-compete agreements affect a variety of margins, but in particular are thought to reduce the frequency of spin-outs from existing firms – that is, firms started by former employees of incumbent firms. For example, Franco and Mitchell (2008) presents a conceptual framework in which lax non-compete enforcement encourages labor mobility and spinouts, albeit at the cost of discouraging firm-specific investments in human capital and also discouraging investments in relationships between customers and suppliers. Empirical work on non-compete agreements has been hindered by a lack of policy variation. Perhaps the best-known empirical work is by Marx, Strumsky, and Fleming (2009), which analyzes a policy reversal in Michigan toward non-compete enforcement in 1985 and finds that average labor mobility dropped substantially, particularly among inventors with firm-specific skills.

Many of the policies discussed in Section 3 – including immigration and non-compete policies but also R&D tax credits and patent box policies – are location-specific policies that may not raise aggregate R&D but rather may simply shift the location of research investments (or researchers) to areas with more generous policy regimes. Lerner (2012) provides an overview of various government programs that have attempted – and usually failed – to encourage economic activity through such place-based policies. Research attempting to isolate evidence for and against various theories of industrial agglomeration – as in the work of Ellison,

[89]See also the excellent paper by Borjas and Doran (2012) on how the large post-1992 influx of Soviet mathematicians affected the publication-based productivity of US mathematicians.

Glaeser, and Kerr (2010) – holds promise for better guiding future policies in this area.

4 Innovation, diffusion, and growth

We have seen that the invention of new technology is an economic problem, the result of directed choices, with incredible importance for growth, industrial organization, and (as we will see in Section 5) inequality. However, the social value of an innovation does not stop with an initial invention, but rather depends also on its *diffusion*.[90] Diffusion refers to the change over time in who produces and uses the invention, and where it is used and produced. We will refer to diffusion as including both technology diffusion, the adoption of a good by new users and in new places, and knowledge diffusion, where ideas spread from one agent to another.

Theoretically, diffusion matters because invention in the ether creates no social value. To be useful, an invention must be either used as an input to production, or as final consumption. Social value thus derives from the invention diffusing to new firms or plants (in the case of a productivity-enhancing invention) or to more or varied final users (in the case of a new consumption good). This is especially relevant as global R&D is highly concentrated in a very small number of countries, and most firms do little R&D themselves: diffusion of technology from those countries and firms is the main source of productivity improvement. Technology can take an exceedingly long time to spread. Indeed, critical inventions like electricity and the computer took decades to become commonplace in industries where both technologies are now fundamental (David, 1990).

This delay has long been noticed by scholars. Indeed, the fact that good ideas do not immediately spread has been studied by sociologists and anthropologists as far back as the late 19th century. Gabriel Tarde, in *Les lois de l'imitation*, argued that even though societal change depends on the diffusion of new inventions, the invention process itself bears "the same relation to imitation as a mountain to a river" (Tarde, 1890). Once invented, they flow from one region to the next according to both the amount of social contact and "logical laws."[91] Chapin (1928) investigates these ideas quantitatively, mapping the cumulative distribution function over time of a number of social innovations such as new forms of city government. It found that an "S-curve" fit these properties well. The famed S-curve is a logistic function of adoption over time with slow initial adoption, then rapid uptake, then a slowdown as a market reaches saturation.

Building on this insight, empirical studies of technology diffusion began with three canonical studies in the mid 20th century. Ryan and Gross (1943) shows that hybrid seed corn was known to many Iowa farmers in the early 1930s, but it was not commonly planted until the end of the decade. In a study of 323 farmers, they find that neighbors attesting to its benefit, rather than salesmen, advertisements, or government extension services, drove adoption. Coleman, Katz, and Menzel (1957) investigates diffusion of "gammanym" (the antibiotic tetracycline) across Chicagoland doctors using prescription records and a survey of doctors about social connections. Again, social ties are relevant for adoption especially early on. Professionally-oriented rather than patient-oriented doctors adopt earlier, S-shaped adoption curves are again evident, and a higher slope of adoption for better-connected doctors is given as evidence of "snowball effects" in learning. Finally, Griliches (1957) revisits the adoption of hybrid corn. The fundamental stylized fact of diffusion is seen in Figure 14, updating his famous chart of hybrid corn adoption by state. Technologies - in this case, hybrid corn - are not adopted immediately. Rather, they are adopted slowly, and the intensive margin of their adoption often appears to follow a logistic curve. Griliches quantitatively decomposes the elements of the S-curve into "origin" (when is a new product introduced), "slope" (how quickly is it adopted), and "ceiling" (at what level does the S-curve plateau with adoption slowing to zero). By looking at S-curves for hybrid corn by region, he shows that economic variables like market density explain earlier introduction, and profit potential explains steeper slope. That is, the problem of diffusion may not be a purely sociological one, but rather a process amenable to economic explanations, and hence economic policy.[92]

[90] For article-length reviews of the economics of diffusion, see Hall (2006), Comin and Mestieri (2014) and Stoneman and Battisti (2010).

[91] See Kinnunen (1996) for further details of the work of Tarde.

[92] The early multidisciplinary literature on diffusion was canonically summarized by Rogers (1962), now in its fifth edition.

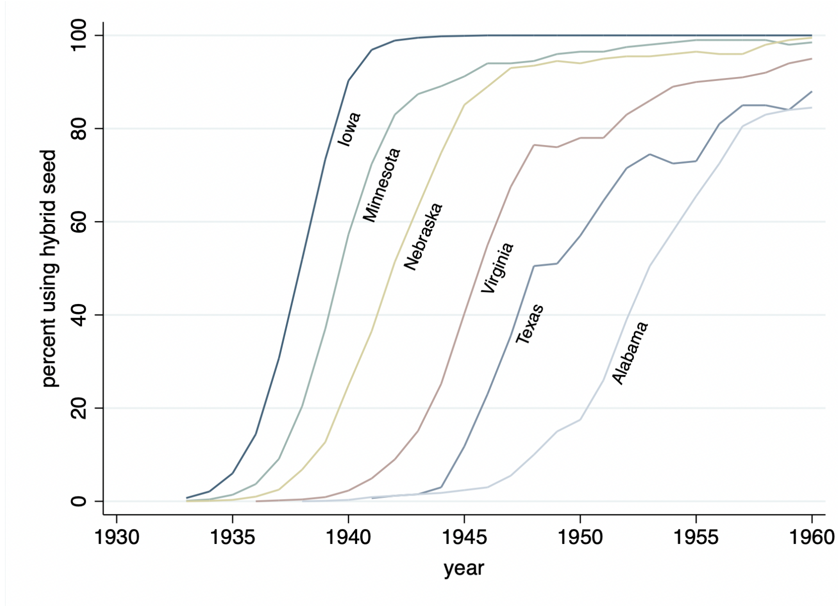

Figure 14: Percentage of total corn acreage planted with hybrid seed

Notes: The y-axis measures the percent of total corn acreage planted with hybrid seed by state, and also nationally. The graph shows the staggered implementation of hybrid corn across the corn belt in the US between 1930 and 1960. Source: Sutch, Libecap, and Steckel (2011), updating data from Griliches (1957).

The time lag in adoption due to diffusion frictions is consequential on micro- and macro-grounds. For example, microeconomically, Skinner and Staiger (2015) finds a three percentage point difference in one-year survival for US heart attack patients depending on whether the hospital they were sent to was a fast or slow adopter of effective treatments like beta blockers, aspirin, and early reperfusion. Differences in acceptance of new technology swamp spending on factor inputs in determining hospital productivity.

On a macro level, diffusion is critical to growth because cross-country differences in per capita output are far too large to be explained by differences in factor inputs (e.g., Klenow and Rodríguez-Clare (1997)). As we saw in the introduction to this chapter, models of endogenous growth (Aghion and Howitt, 1992; Romer, 1990) rely on two fundamental factors: invented technology is non-rival, and inventors must capture at least some of the social surplus of what they invent. Drawing on micro-studies of diffusion, Eaton and Kortum (1999) gives a tractable structural model of innovation diffusion across regions. Foreign inventions raise domestic TFP as a function of the rate at which that knowledge arrives. This growth effect can either come from the use of new inventions as inputs into final goods in the destination country, or the idea in those inventions as the source of sequential inventions. Factors which drive the diffusion rate are, in this model, the primary explanation of differences in the wealth of nations. The study of innovation is therefore incomplete if it stops with invention and initial commercialization, without considering the factors that lead to diffusion and broad adoption.

4.1 Measuring diffusion

Before turning to the theory and empirics of *why* diffusion may be too slow, let us examine what data exists to study this question at all. Going back to Ryan and Gross (1943), most diffusion studies rely on hand-collected data about specific technologies. For instance, the classic study of industry growth following a canonical invention in Gort and Klepper (1982) draws on panel data from 46 industries, as well as custom constructions of "major" and "minor" inventions. In general, government and private sector datasets tend to have very limited data on the diffusion of particular technologies, and tracing the source of that diffusion is even more difficult. Nonetheless, there are four classes of nonproprietary data which have proved useful:

CHAT, scanner data, census data, and patent data.

In terms of pure data on technology adoption over time and space, including the intensive margin of adoption in each location, the Cross-Country Historical Adoption of Technology (CHAT) dataset introduced in Comin and Hobijn (2010) is the most extensive. This dataset traces the intensity of adoption of over 100 technologies in 161 countries since 1800. Having both adoption and the intensity of adoption by country is useful: though the airplane was available in both China and the US in 1960, the number of flights (as measured by passenger-miles) was orders of magnitude higher in the US. Comin and Mestieri (2018) argues that the intensive margin of diffusion, rather than the extensive margin, explains why incomes diverged between countries even as technology "arrived everywhere." Going back further in time, the Primitive Technology dataset attempts a similar computation for 1000 BC, 0 AD, and 1500 AD (Comin, Easterly, and Gong, 2010).

Scanner datasets, generally limited to a single country, contain panel information on precisely which products (as measured by, for instance, UPC or Universal Product Codes) are available for sale when and where.[93] While some datasets are limited to specific industries (for instance, the IRI Academic Dataset on foodstuffs), the standard cross-industry reference in marketing and industrial organization is the Nielsen Retail Scanner Dataset. This dataset contains weekly data for millions of consumer-oriented UPCs from a sample of tens of thousands of stores going back to 2006. Though these data are less commonly used by innovation scholars (Argente, Baslandze, Hanley, and Moreira (2020) is an exception), there is tremendous opportunity in linking the diffusion data in these UPCs to other covariates about the invention in patent or hand-collected data.

A third source of data is large-scale surveys, and in particular census surveys. For certain classes of consumer goods (e.g., automobile or air conditioner ownership), standard population surveys in many countries have long tracked these data. For more specific technology, especially those used as a process input by firms, we know much less. Nonetheless, two specific surveys appear particularly promising for studies of diffusion to firms. The US Census Annual Business Survey, whose first wave was in 2018, asked 850,000 nationally representative businesses questions about their adoption of technologies including cloud computing, machine learning, radio-frequency identification, robotics, and more. As one illustration, Figure 15 documents adoption of a number of frontier technologies, particularly related to cloud computing, in these data. For investigating global firm adoption, the World Management Survey, which derives from Bloom and Van Reenen (2007), asks managers in a growing number of countries questions about their adoption of process technologies.[94]

Finally, as with many other questions in innovation, the "paper trail of knowledge" in patents has been a useful dataset for studying diffusion.[95] Jaffe, Trajtenberg, and Henderson (1993) canonically shows that patents are two to six times as likely to be cited by a future inventor in the same metropolitan area as a random patent with similar technology and total citations. This localization effect suggests that the ideas in a patent diffuse over space only with time. Nonetheless, one may quibble that this limited diffusion of ideas is driven by imperfect matching between control and focal patents. Thompson and Fox-Kean (2005) takes advantage of a 2001 policy change where the USPTO began denoting which citations were added by examiners to allay that concern. Citations from inventors are much more likely to be from near the inventor's location than examiner-added citations, again suggesting the importance of geography in diffusion. Arora, Belenzon, and Lee (2018) offers a skeptical note when considering "citation reversals" where a patent cites another patent which is newer rather than older. That is, a patent with priority year 2000 cites, in its 2008 application, a 2004 priority year application published in 2006. This cannot reflect true temporal knowledge spillover, by definition. However, these citation reversals have almost exactly the same geographic decay as

[93] For pharmaceuticals, similar UPC code-level data on revenues and quantities by country is available from the firm IMS Health, as used by Costinot, Donaldson, Kyle, and Williams (2019).

[94] A similar US survey, the Management and Organizational Practices Survey, finds very slow diffusion of management practices even within a single firm, especially in non-competitive sectors (Bloom et al., 2019).

[95] See also Adams and Clemmons (2013) measuring lags between citing and cited scientific papers.

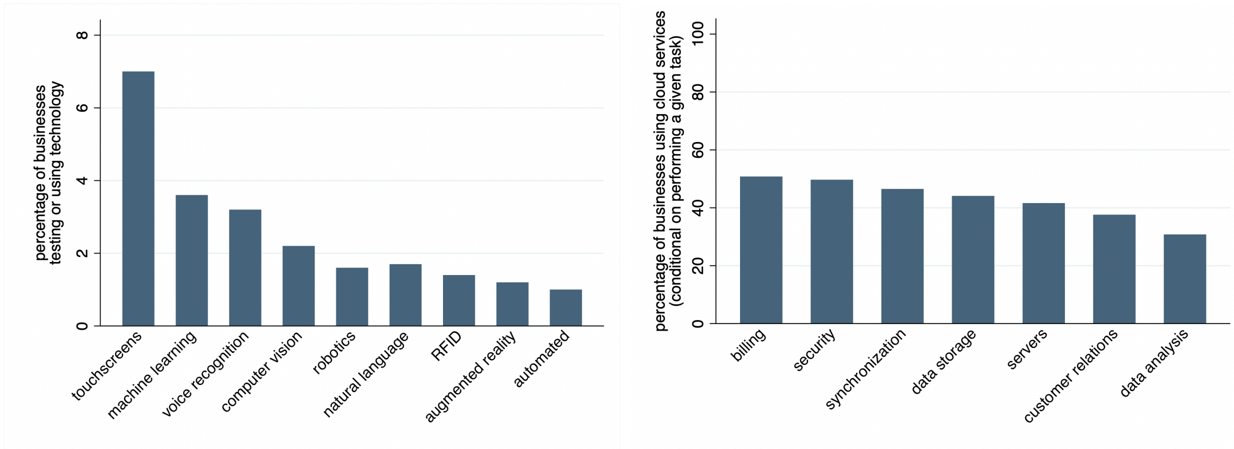

(a) Adoption or trial of advanced technologies

(b) Adoption of cloud services

Figure 15: Heterogeneity in adoption of technologies

Notes: The left panel shows adoption or trial of advanced technologies as reported in the 2018 Annual Business Survey. The right panel shows the fraction of businesses who use cloud services - those performed by a remote server using an internet connection - conditional on performing each task in their business. These are both weighted to match all US firms on the basis of responses from over 570,000 businesses. Source: McElheran et al. (2020).

non-reversal citations.

A particularly intriguing technique to measure diffusion using patents is to draw on the text itself using machine learning, rather than relying only on the database of citations. Myers and Lanahan (2021) investigates spillovers from SBIR grants targeted at particular technical problems, taking advantage of heterogeneity in state matching funds to account for selection on the technological areas being funded. Rather than looking at backward citations to patents received by the grant winner, they look at the cosine similarity between the text in the SBIR call for proposals and the full corpus of US patents in order to infer precisely which technology classes the grant was targeting. They find that up to four patents overall are generated for each patent induced by the grant recipient, that the citation network misses up to half of these spillovers, and that spillovers are in fact more likely to be geographically proximate.

The intensive margin of new technology adoption as an input to other inventions, in addition to geographic spread, can also be traced by patents. For example, Pezzoni, Veugelers, and Visentin (2019) considers sets of patent classes combined on a single patent. For example, in 1985, "gene isolation" and "injection of material into animals" were joint classes for the first time in the oncomouse patent. They show that the cumulative number of patents using new class combinations grows over time according to S-curves. Patent class combinations which are more similar are used by other inventors very quickly, but the total number of patents ever using that combination is low, compared to more unusual class combinations. Combining patent data with information about invention and external data on adoption by end users or non-patenting intermediate users is a promising future path. Further, the text of patents themselves often contains information written by inventors about the precursors of their invention. Standardized datasets of scientific references in patent text now exist (Bryan, Ozcan, and Sampat, 2020; Marx and Fuegi, 2020), but more complex uses of natural language data to study technology diffusion are underexploited.

4.2 Theories of diffusion

The mid-century studies we have seen suggest two factors affecting diffusion. First, information asymmetries mean that social networks smooth diffusion, and second, that adoption is costly in both a real sense and in terms of the option value of waiting to learn more. More recent literature focuses on two further classes of explanation: heterogeneity in adoption costs or benefits, and "goldbricking" incentives among labor to prevent adoption. Let us tackle these in turn.

Asymmetric information, and hence the role for social contacts, as an explanation for slow diffusion goes back to sociological studies of ideas spread in the nineteenth century. As a result, firms introducing new products often make them similar to existing products in the same category in order to smooth this learning. For instance, Hargadon and Douglas (2001) shows that Edison deliberately made the electric light as similar to existing gas technology as possible, with a centralized power source rather than in-building generators, metered use (despite the lack of an effective meter), wattage similar to existing gas flames, and buried wires despite the high cost of fixing short circuits.

When the benefits of new technology are hard to observe, social learning – where agents learn from friends and neighbors about the benefits of a new technology – can drive adoption. Empirically documenting evidence of social learning is a challenging problem given that social contacts likely share many covariates in common. Conley and Udry (2010) takes advantage of the lengthy period between planting and the revelation of information about crop output to identify social learning as an important driver of the diffusion of fertilizer-heavy pineapple crops in Ghana. Because the link between action and outcome is stochastic, and because social learners sometimes only observe the outcomes of a small number of their social contacts, diffusion may be limited by "bad luck." Adhvaryu (2014) considers a new malaria medicine in Tanzania which was rolled out quasi-randomly. Misdiagnosis of malaria and fever slows social learning, and lowers future adoption of this "technology" in villages where misdiagnoses as measured by later blood tests was high. Broadly, subsidies for tools that help social learning, like better diagnostics, may be economically valuable. Information asymmetries can also be solved with non-social information: Arrow, Bilir, and Sorensen (2020) finds that doctors with access to a pharmaceutical reference database are much more likely to prescribe cheaper generic drugs instead of their branded equivalent.

A theoretical caveat on social learning is important to keep in mind given the interdisciplinary nature of diffusion studies. The implicit models used in sociology, marketing, and economics are not the same even if the descriptive language is similar. In particular, sociological and marketing models implicitly assume that diffusion requires social contact (a "contagion" model) or is based on wanting to conform to the choices of neighbors (a "social influence" model). Social learning instead involves rational updating of beliefs about a technology on the basis of observing the choices and outcomes of friends. Young (2009) shows theoretically that contagion, social influence, and social learning models can be distinguished by the nature of the adoption curve.

The second factor affecting diffusion, the costs of adoption or information transmission, relies on the idea that many inventions are tacit rather than codified. In order to be exploited, they must be explained, and doing so is not easy. Even within a multinational firm, Teece (1977) shows that the costs of transferring knowledge about production made up an average of 20 percent of the costs of moving production to a new plant in a multinational firm. Glitz and Meyersson (2020) suggests quantitatively that active corporate espionage was required to keep East Germany from falling further behind the productivity frontier in the West. These findings are both consistent with Arrow's dictum that costly communication of tacit knowledge makes technology transfer challenging even when incentives are properly aligned (Arrow, 1969). Regions with a common language produced disproportionate numbers of important inventions historically, particularly those requiring cooperation among many inventors (Dudley, 2017). Even the literal construction of roads connecting two cities, in a study comparing built with planned but unbuilt infrastructure, causes an increase in patents which draw on local knowledge (Agrawal, Galasso, and Oettl, 2017). Understudied are factors which reduce that adoption or information cost, smoothing technology adoption across firms, consumers, or regions.

Indeed, Schumpeter (1934) places this factor front and center: in addition to creating new products or processes, the "opening of a new market...into which the particular branch of manufacture of the country in question has not previously entered" is itself a costly form of innovation. Therefore, supplier incentives matter. This is especially true if the inventor has a unique ability to supply a product, because of IP protection or unique firm resources. For instance, Kyle (2007) and Cockburn, Lanjouw, and Schankerman (2016) find that new drugs diffuse more quickly to countries with stronger patent rights, less price regulation, and more limited market regulation. These are not minor effects, as many new drugs only become widely available a decade or more after their initial approval.[96]

The costs of adoption are not just real costs, but also option value. The decision to adopt a technology *today* must be weighed not just against nonadoption, but against adopting tomorrow. When adoption incurs a fixed cost, and information about the benefit of adoption is revealed over time, there is a value to waiting. Worse, the decision to wait under social learning also imposes an externality to future adopters in one's social network. Farzin, Huisman, and Kort (1998) shows theoretically that diffusion is slower than the optimum for a wide variety of sources of uncertainty about a new technology, a result strengthened further by a technical correction in Doraszelski (2001). The option value of potentially free-riding on external innovation in the future distorts the rate of invention today: even technological leaders realize that the value of "falling behind" is not so bad if frontier technology developed by others will quickly diffuse (Benhabib, Perla, and Tonetti, 2014). A long debate in industrial organization concerns the extent to which irreversible adoption decisions combined with the option value of waiting causes inefficient technological direction. David (1985), Cowan (1990), Farrell and Saloner (1986) and Arthur (1989) discuss the basic possibility of "lock-in." In terms of the empirical importance of lock-in, see Liebowitz and Margolis (1995) for an argument against, and Farrell and Klemperer (2007) in the (partial) affirmative.

The combination of option value to adopting and social learning has been studied by development economists investigating why some highly productive inventions do not diffuse widely in developing regions. Foster and Rosenzweig (1995) finds that high-yield seeds in India are underused because farmers must experiment to learn their effectiveness, that they can learn from others in their village, and that they underexperiment due to this positive spillover. Bandiera and Rasul (2006) investigates sunflower seed adoption in Mozambique, again finding that adoption by contacts of the same religion, or by friends and family, improves adoption probability. Given these results, the effect of ethnic or religious cleavages on information transfer, or the benefit of technology solving this information problem, are questions relevant to addressing global poverty.[97]

One solution to inefficiencies caused by the option value of delaying adoption are standard-setting organizations (SSOs). SSOs attempt to avoid wasteful innovation and marketing of technologies that will not be adopted; examples include intergovernmental bodies like the International Air Transport Association and private associations such as the Internet Engineering Task Force. Simcoe (2012) develops a theoretical model of voluntary SSOs, arguing that when the transaction costs of side payments or the private benefits of pushing a preferred technology over the standard increase, standards are less likely to be rapidly accepted. He shows precisely this delay empirically in the early days of the commercial internet. Since SSOs are voluntary, their internal organization and development are particularly important for efficiency. Two aspects have drawn particular attention: how to prevent "standard-essential" patent holders from demanding high license fees once the standard is established (Lerner and Tirole, 2015) and the implication of "forum-shopping" among different SSOs (e.g., Lerner and Tirole (2006)). To the extent that SSOs not only develop standards for interoperability and coordinate complements over time, but also involve reciprocal patent licenses, they may also face some of the efficiency concerns that arise with patent pools (Hagiu and Yoffie, 2013).

A third factor driving diffusion is that potential adopters have heterogeneous preferences or benefits from the technology. For instance, the Kenyan government found maize yield increases from fertilizer and hybrid seeds of 40 to 100 percent in experimental model farms, yet only 60 percent of small farmers adopted these technologies. Duflo, Kremer, and Robinson (2008) induces small-scale farmers to use the government-recommended amount of fertilizer on part of their plot and otherwise to farm as they normally do. Simply

[96] On pharmaceutical regulation, see also the work of Henry, Loseto, and Ottaviani (2021) on how regulators should handle private information about benefits and safety characteristics of new drugs.

[97] See Foster and Rosenzweig (2010) for a review of the literature on the microeconomics of diffusion and its link to development.

using the recommended fertilizer level has a negative expected return, hence would not be adopted in the absence of other organizational or technological changes. Heterogeneity also means that some technologies will not be widely used until complementary technologies are invented (Bresnahan and Trajtenberg, 1995). Gross (2018) examines the spread of the tractor across the US in the early twentieth century. Fixed-tread tractors were widely used in the Wheat Belt in the 1920s and early 1930s. Only with the arrival of a general-purpose tractor, and complementary technology that permitted common cotton and corn farm tasks to be performed, did the tractor spread widely to Midwestern farms. That is, the tractor was not adopted in the 1920s by many farms because it did not have the scope of use cases, and the complementary technology from third parties, to make it worth buying for corn and cotton farmers. The role of complementary technology in adoption is particularly important when multiple agents must adopt simultaneously for a technology to be valuable. Basker and Simcoe (2021) examines adoption of UPC codes, finding that manufacturers are more likely to adopt UPC when other manufacturers sharing the same retail partner have adopted them.

Heterogeneous preferences or benefits from a new technology mean that diffusion increases as the cost of technology falls. Consider the role of cost in the adoption of the spinning jenny (Allen, 2009). The spinning jenny – a labor-saving invention in textiles which played a fundamental role in the Industrial Revolution in England – did not diffuse widely in France or India in the late 18th century. Why? Allen estimates that wages were not high enough in France or India for the spinning jenny to be cost effective until the relative price of the device compared to wages fell. Since prices of new inventions often fall due to scale economies or learning-by-doing (Arrow, 1962b), diffusion lags may result simply from the fact that early adopters have higher value for the invention. This can be true even for inventions which are used in production, as in the case of the spinning jenny.

An important reason why adoption costs for new innovations differ is the highly local nature of knowledge spillovers. Audretsch and Feldman (1996) investigates the geographic specialization of new innovations vis-a-vis production (in terms of employment or value-added) across US states. Particularly in knowledge-intensive industries like computers or pharmaceuticals, innovation production is much more geographically specialized than final goods production. The same importance of very local spillovers is evident in the geography of citations between patents (Jaffe, Trajtenberg, and Henderson, 1993) and the concentrated location of high-growth new firms (Guzman and Stern, 2020). If knowledge flows are local, there are positive externalities from being near other innovators. Local technology will diffuse more quickly due to lower adoption and search costs. These externalities have important policy implications: for example, *local* tax incentives to attract innovative firms internalize local but not global economic benefits, potentially pulling firms away from highly productive agglomerations (Slattery, 2020).

Finally, diffusion may be limited because some agents benefit from preventing the new technology from being used. That workers can take actions to conceal the effectiveness of new technology, particularly labor-saving technology, is called *goldbricking* (Roy, 1952). For instance, workers during the Industrial Revolution resisted the adoption of the self-acting mule which was thought to reduce labor demand (Lazonick, 1979). When compensation structures are sticky, even labor-enhancing inventions may be rejected. Atkin, Chaudhry, Chaudry, Khandelwal, and Verhoogen (2017) shows that soccer-ball producers in Pakistan do not adopt a highly-efficient new method of preparing soccer balls because piece rate workers are not compensated for avoiding waste in production, but do lose wages when their production rate slows as they try to incorporate the new technique. To the extent that many organizational practices are necessarily sticky, perhaps due to relational contracts, useful technologies may not diffuse rapidly (e.g., Gibbons and Henderson (2011)).

These factors limiting diffusion have substantial consequences across many areas of economics, including health, trade, and antitrust, among many others. Agha and Molitor (2018) finds that when cancer treatment clinical trials are led by a local scientist, patients in that region are 36% more likely to be given the drug in the two years after FDA approval: information frictions to diffusion hold for research-oriented physicians in the same way they do for Ghanaian pineapple farmers. Allen, Bilir, Chen, and Tonetti (2019) examines prescriptions of cholesterol treatments among a sample of 50,000 doctors, finding strong evidence of learning over time as a function of geographic and medical school cohort similarity with doctors who learn about more effective treatments early. Their model suggests that targeting the initial beliefs of doctors is much less effective in diffusing accurate information than is strengthening information flow over time in highly-connected parts of the network. As is well-known in the network theory literature (e.g., Akbarpour, Malladi, and Saberi (2020)), the optimal way to speed diffusion across a network by targeting agents or their connections is often quite subtle. At the level of country adoption rather than individual adoption, Kyle (2006) shows that many new drugs are introduced to some countries only with substantial delay, or never launched at all. When the inventing firm has experience in a country, or in neighboring or culturally-similar countries, they are substantially more likely to introduce the drug. This suggests that some technologies do not diffuse because of a combination of diffusion frictions with limits on the ability to license to other firms better able to diffuse the product to certain locations.[98]

In standard trade models, the gains from free trade are often thought to be counterintuitively low. Dynamic gains from trade, which incorporate the effect of openness on the firm productivity distribution, either from selection or from access to ideas, are a potential solution. Buera and Oberfield (2020) gives a tractable model of diffusion where the critical parameter is how important global ideas are to local productivity. The benefits of trade to growth are highest when this parameter is in an intermediate range: foreign ideas are important to productivity, but not so important that you can learn everything useful by importing a small variety of foreign products. The nature of frictions to diffusion is thus critical for welfare analysis of trade policy.

A common policy intended to speed diffusion when trade alone is insufficient is a technology transfer requirement. Whether using patent linkages (Jaffe, Trajtenberg, and Henderson, 1993) or direct measures of industry-level total factor productivity (e.g., Keller (2002)), inventions appear to diffuse slowly over space. Policies that increase technology transfer, either through IP rules that incentivize this transfer or direct technology transfer requirements for producers, may therefore be useful for recipient nations. Scotchmer (2004) argues theoretically that reciprocity in treatment of foreigners in IP protection improves welfare by solving the free-riding problem of low innovation countries. However, forced technology transfer also acts as a tax on initial innovators by shrinking the gap in productivity between innovators and beneficiaries of transfer rules. The global welfare effects of these policies are therefore subtle, and require further investigation.[99] Beyond mandated technology transfer requirements, governments and industry groups have also long attempted to directly subsidize the diffusion of new ideas, such as in 19th century Britain with the Society for the Diffusion of Useful Knowledge. More recently, and more globally, policies like agricultural extension services (e.g., Birkhaeuser, Evenson, and Feder (1991) and Kantor and Whalley (2019)) or deliberate transfer of productivity advice (Giorcelli, 2019) are frequent components of domestic and foreign policy. Giorcelli (2019) examines a Marshall Plan program to train Italian managers and provide frontier technology. The program's budget was unexpectedly cut, leading to exogenous geographical variation in participation. The transfer of management knowledge, in particular, led to increased productivity and a higher probability of exporting even fifteen years after the program ended, while pure technology transfer showed more temporally limited benefits.

In contrast to forced technology transfer, some countries have explicitly *taxed* foreign technology. The goal is both to promote domestic innovation, and to ensure that innovation is more "appropriate" for the skill mixture in the domestic labor force. de Souza (2020) combines a directed technical change model with a unique dataset from Brazil where international technology licensing must be recorded with the government. A 2001 law taxing foreign technology and subsidizing domestic innovation led firms that previously used

[98] Indeed, when diffusion is important, patents or other sources of market power may incentivize marketing, a form of non-market competition. Lakdawalla and Philipson (2012) finds that slower spread of new drugs due to decreased marketing means the consumer surplus from drug patent expiration is less than the pure price change would imply. This argument is in line with the prospect theory of patents (Kitch, 1977), which argues that patents not only incentivize invention, but also the optimal further exploitation of the patented idea, similar to how real property rights encourage exploitation of owned land.

[99] See Branstetter, Glennon, and Jensen (2019) for a summary of issues related to US-China technology transfer rules.

technology from abroad to increase patenting, hire lower skilled workers, and decrease production. Broadly, the global level of technology is of sufficiently high quality that skill mismatch alone does not mitigate against its use.

A caveat to the effects of technology transfer rules is the idea of absorptive capacity, particularly for inventions that lift firm productivity (Cohen and Levinthal, 1989). Absorptive capacity means that, both across and within firms, the ability to absorb invention from the outside depends on what the recipient knows.[100] That is, a plant manager with her own R&D division is better able to keep abreast of, and evaluate, potentially useful inventions. Bilir and Morales (2020) estimates that 20% of the productivity benefits of innovation in the US by the median multinational firm accrues to foreign affiliates, and that there is strong complementarity between those affiliates R&D and this productivity gain. Coe, Helpman, and Hoffmaister (2009) likewise finds that foreign R&D stocks benefit domestic TFP especially when the recipient nation is highly-educated and has strong IP protection. That is, the gains from multinational production do not diffuse for free, but rather require active investments in the capacity to learn, and suitable institutions for encouraging technology transfer by recipient plants, firms, or countries.[101] The need for absorptive capacity may therefore partially counteract the harms to R&D incentives caused by technology transfer requirements. The precise balance of these effects empirically is an open question.

The stylized facts on diffusion also suggest an intriguing new problem for antitrust. Slow diffusion means that productivity differences between firms can remain for long periods of time even in the absence of anti-competitive behavior. Syverson (2004) documents that the 90th percentile firm in a given 4-digit SIC industry is almost twice as productive as the 10th percentile firm, a figure which is even larger in a global context (Hsieh and Klenow, 2009). If technology diffusion requires active effort, policies which increase competition may change the incentive of incumbents to adopt frontier technology or force low-productivity firms out of business if they do not adopt those techniques. Schmitz (2005) finds that competition with Brazilian mines forced American and Canadian iron ore mines to adopt organizational practices and labor contracts which doubled productivity in the early 1980s. Bloom, Draca, and Van Reenen (2016) finds that lower tariffs for Chinese producers forced European manufacturers to either adopt productivity-enhancing technology like frontier IP or else to exit the industry. Studies in this area tend to focus on aggregate TFP or labor productivity measures rather than direct technology or organizational practice diffusion. Direct studies of how changes in competition affect the rate at which new technology spreads within and across firms remain limited.

More broadly, the theoretical nature of externalities in diffusion - inclusive of the many factors why technology diffuses slowly discussed in this section - is underdeveloped. Empirically, there are very few credible estimates of the social return to policies which speed diffusion in the aggregate, with social network-based technology adoption a notable exception. This is true even though specific inventions have enormous social returns once adopted: the printing press doubled the relative wage of the university-educated in Europe while lowering the price of books in regions where it arrives by a factor of 100 (Dittmar, 2011). Research understanding where equilibrium investment in marketing and diffusing new products or processes is inefficient, and which public policies can ameliorate this inefficiency, is particularly needed.

5 Innovation and inequality

Our focus in Section 3 was on assessing how government policies which attempt to shape market incentives – either through reducing the costs of research, or increasing the expected revenues from research – affect research investment. Some of the papers attempt to move beyond research investments to evaluate outcomes more directly linked to social welfare, but even then the focus is on efficiency. That is, when will research be

[100] In the context of individuals, Nelson and Phelps (1966) also argues that human capital is relevant for the adoption of outside technology.

[101] Keller (2004) provides a detailed evaluation of the challenges in estimating the causal benefits of technology transfer.

under-provided by the private market, and how well do different policy levers rectify inefficiency?

A different but critically important question is how these policy levers affect *inequality*.[102] In recent decades, the US as well as many other developed countries have seen sharp increases in income inequality (Piketty and Saez, 2003), and much attention has been focused on the extent to which market power and other economic structures have contributed to this trend. From an innovation policy perspective, several questions are relevant. Has innovation contributed to the observed rise in inequality? How do policies like the patent system – which by construction influence how the returns to new inventions are divided among inventors, firms and consumers – affect inequality? At a societal level, does inequality in opportunity lead us to lose potential innovators and, in turn, potential inventions? The economic literature on these questions is still very much in its infancy, but we here highlight some relevant work with an eye towards highlighting areas with particular promise for future work.

5.1 Does innovation increase inequality?

Theoretically, innovation should affect inequality across a variety of models. One example is the Schumpeterian growth model of Aghion, Akcigit, Bergeaud, Blundell, and Hemous (2019), in which growth results from quality-improving innovations by either incumbents or potential entrants. Easier quality improvements by either the incumbent or entrants allow higher markups, and hence increase inequality. Their model also predicts that entrants' innovation increases social mobility, which suggests that fostering entrant innovation results in growth that is more inclusive. On the other hand, Jones and Kim (2018) finds a more tenuous link between top inequality and innovation. In their model, incumbents can improve existing technology only over time, and with some probability are replaced by a disruptive technology. As it becomes easier for incumbents to build on existing ideas, such as through access to more markets via globalization, inequality increases. In the short run, there is a clear link between easier innovation and top inequality. However, when incumbents are better able to improve their existing products, variance in the payoff to R&D by entrants goes up, with researchers producing either quick-growing disruptions or failed attempts to disrupt. Since inventors are risk-averse, fewer are willing to accept that gamble, hence both the rate of creative destruction through new innovations and the overall growth rate of the economy fall. In the long-run, therefore, sclerotic economies with powerful incumbents have high inequality and low growth, while innovative economies see *decreases* in top inequality and faster growth.

Empirically, the question of whether innovation increases inequality is one that scholars have investigated at both a macro-level and at a micro-level. At a macro-level, Aghion, Akcigit, Bergeaud, Blundell, and Hemous (2019) starts by documenting descriptively that changes in innovation – as measured by log patent citations per capita at the state level in the first five years after patent applications are filed – are correlated over time with state-level increases in the share of income held by the top 1%. The paper then develops an instrumental variables (IV) strategy to test for a causal channel running from innovation to inequality, arguing that the composition of the US Senate Committee on Appropriations affects the allocation of earmarks across states, which in turn affects patenting and innovation at the state level. Their IV estimates leveraging this idea suggest that an increase of 1% in patents increases the top 1% income share by 0.2% – consistent with an economically meaningful and statistically significant causal relationship running from innovation to inequality.[103]

While they are drawn from a diverse set of distinct literatures, we are aware of three papers offering more micro-level perspectives on how innovation can affect inequality. Cutler, Meara, and Richards-Shubik (2012) starts with the thought experiment of a social planner wishing to maximize social welfare by minimizing mortality rates subject to a fixed budget constraint, as might be the case for the US National Institutes of Health (NIH). Equalizing the expected marginal benefit of research across diseases gives three predictions.

[102] Bloom, Van Reenen, and Williams (2019) explicitly includes an assessment of the likely effects of different innovation policies on inequality in their review; see their Table 2.

[103] Also related is the work of Berkes and Gaetani (2019), which investigates how local innovation affects income segregation.

First, initial death rates should be positively correlated with subsequent research efforts. Second, innovation should – assuming innovation is productive – be associated with more rapid mortality declines. And third, induced innovation will increase mortality disparities between minorities and the majority group if minorities and the majority group differ in their shares of deaths from "leading" causes. It is this last prediction that is most of interest: if the NIH has what seems ex ante like a very reasonable objective function of targeting subsidies at the diseases that are most prevalent in the population as a whole, but if e.g. Blacks suffer from different causes of deaths than whites on average, then NIH funding could *increase* racial disparities in mortality outcomes. Importantly, this prediction arises not because of any form of bias but rather is just a consequence of social welfare maximization in the set up of this model. Cutler, Meara, and Richards-Shubik empirically assess these predictions in the context of trends in infant mortality (that is, deaths during the first year of life) in the US from 1950-2007.

Consistent with the first two predictions of their model, population-level infant mortality rates (deaths per 1,000 live births) in 1983-85 are positively associated with levels of NIH funding over the subsequent 15 years, and diseases which received more NIH support saw larger decreases in mortality rates over a later time period (from 1996-1998). To analyze their third prediction, the authors start by noting that over this period, the US had around five times as many white births as Black births, and the overall leading causes of death are those from which whites suffer from relatively more. Hence, when progress was made on those causes of death that were most prevalent in the population as a whole – such as respiratory distress syndrome, for which the drug surfactant was an important new treatment diffusing over this period – the authors present simulations suggesting that racial gaps in birth weight-specific mortality widened over time as a result of this progress. That is, public research support responding to aggregate medical needs improved aggregate outcomes, but a disproportionate share of those benefits accrued to white infants.

While the motivation for his study is quite different, Jaravel (2019) presents an analysis quite similar in spirit to Cutler, Meara, and Richards-Shubik (2012). Jaravel starts by contrasting two narratives about who benefits from innovation. A first – more classic – narrative focuses on the idea that everyone benefits from innovation via a product cycle. Innovation may initially be aimed at the rich, who buy first, but a trickle-down process eventually brings products to the mass market. A second, alternative narrative is that rising income inequality implies that demand will grow faster for "premium" products (e.g. organic products) that are disproportionately consumed by richer households. In this latter case, endogenous entry of new products may lead to an increase in product variety and reduced prices in premium categories predominantly consumed by high-income households, generating disproportionate benefits to that group. Jaravel focuses on empirically assessing the relevance of this second narrative using scanner data. He documents that higher-income households experience faster increases in product variety and lower inflation in the US retail sector from 2004-2015, and then presents a back-of-the-envelope calculation suggesting that this channel may account for a large fraction (over 50%) of observed inflation inequality.[104]

Finally, the third micro-level analysis of innovation and inequality is drawn from the labor economics literature on rent-sharing.[105] A long literature in economics has documented evidence that some firms – in particular, more profitable firms – appear to pay higher wages than others, even when (measured) labor quality and occupation are held constant. In an important contribution, Van Reenen (1996) proposes using the quasirents earned by firms developing technological innovations as a source of quasi-experimental variation in firm rents, and then asking whether those (instrumented) firm rents are passed through to workers in the

[104] In a closely related paper, Jaravel (2018) explores the implications of this mechanism for the food stamps program, a US transfer program which provides benefits to low-income households to help pay for the cost of food. Specifically, he empirically documents that food stamp eligible households experienced relatively lower inflation and relatively higher product variety in states with a larger increase in the take-up rate for food stamps, compared to food stamp ineligible households.

[105] Economic theory has, going back at least to David Ricardo, also considered the question of who benefits from innovation. Ricardo argued, in a model where land is a fixed factor and human population expands in a Malthusian manner, that innovation could both reduce overall wages and overall economic surplus inclusive of wages. See Samuelson (1989) for a modern statement of this result, and for a discussion of precisely when innovations which expand the production possibilities frontier nonetheless reduce total surplus paid to a given factor like labor.

form of higher wages. Unfortunately, the only wage data he had available was firm-year aggregates, which limited his ability to test for a causal relationship between innovation-generated rents and wages given that compositional changes in who was employed at the firm could also have generated observed changes in wages.

Kline, Petkova, Williams, and Zidar (2019) overcomes this limitation by developing a new linkage between USPTO administrative data and US Treasury firm and worker tax filings. Variation in USPTO initial allowance decisions identifies the causal effects of patent allowances on firm and worker outcomes. Consistent with the estimates in Van Reenen (1996), Kline et al. finds that patent allowances raise average earnings at the firm level. However, the micro-data enable them to observe that patent allowances also exacerbate within-firm inequality on a variety of margins. 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. Earnings of male employees rise strongly in response to a patent allowance, whereas earnings of female employees respond less. Such findings suggest that firm shocks play an important role in generating earnings inequality both across and within firms.

5.2 Is inequality causing society to lose potential innovators?

While innovation can affect inequality, the level of resource or opportunity inequality in society also might affect innovation. Much recent attention and concern has focused on the question of whether inequality is causing society to lose potential inventions or innovators. That is, does inequality result in "missed opportunities" where individuals who could contribute socially valuable ideas fail to have the opportunity to realize their potential?

A natural – even if descriptive – starting point is to document how the likelihood of innovating differs across demographic groups. However, in practice most large administrative data sets recording measures of innovative activity record few or no demographic variables. Records of patent applications filed with the US Patent and Trademark Office (USPTO), for example, do not record individual characteristics such as race or gender.[106] In the absence of directly-measured data on demographics such as race and gender, researchers have undertaken a variety of efforts to indirectly infer these variables.

Several papers, such as Jensen, Kovács, and Sorenson (2018), assign probable gender of inventors by matching the first names of inventors to the gender distribution of those names in data such as the US Social Security Administration data.[107] Cook (2014) identifies African American inventors by hand-linking patent records to surveys Henry E. Baker conducted on behalf of the US Patent Office in 1900 and 1913,[108] by matching patent records to US census data, and by curating lists of African American scientists, engineers, and medical doctors based on qualitative sources such as obituaries in local newspapers and published biographies. Cook (2014) leverages these data to document that the gap in patenting between African Americans and whites is larger during periods of ethnic and political violence.[109]

[106] Of course, smaller survey datasets do frequently measure gender directly. For example, Whittington (2009) analyzes gender differences in patenting and publishing in the National Science Foundation (NSF)'s Survey of Doctoral Recipients, and Hunt, Garant, Herman, and Munroe (2013) analyzes gender differences in patenting in the National Survey of College Graduates. Recent policy proposals such as the IDEA Act of July 2019 have encouraged the USPTO to collect demographic data from applicants, but that reform has not been implemented at least as of this writing.

[107] Jensen, Kovács, and Sorenson (2018) applies these data to document differences in patenting across men and women. Similar methods are used by Kolev, Fuentes-Medel, and Murray (2020) to document gender differences in grant writing and grant outcomes, by Ewens and Townsend (2020) and Gompers and Wang (2017) to document gender differences in firm financing, by Marschke, Nunez, Weinberg, and Yu (2018) to document gender differences in last authorships in the biomedical sciences, and by Ding, Murray, and Stuart (2006) to document gender differences in academic life sciences patenting. One validation exercise of this type of methodology is presented in Appendix I of Toole, Saksena, deGrazia, Black, Lissoni, Miguelez, and Tarasconi (2020), which compares probable gender with actual gender as measured for patent *examiners* in internal USPTO human resources data.

[108] Baker sent surveys to 9,000 of the 12,000 patent attorneys and agents in the US at the time, asking if they had African American clients or if they knew any African American patentees. See Cook (2014) Footnote 18 and Appendix I for more details.

[109] Cook and Kongcharoen (2010) combines these two approaches – assigning female names probabilistically and leveraging Cook's earlier work to identify African American inventors – in order to document gender and racial gaps in commercialization activity, as measured by assignment of patents to a corporation, university, organization, or anyone other than oneself as of the date of patent issue.

A trio of recent papers have constructed data enabling a broader analysis of the demographic factors correlated with becoming an inventor: Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019c) for modern US data (linking patent records with administrative US tax records), Akcigit, Grigsby, and Nicholas (2017b) for historical US data (as in Cook (2014) and Sarada, Andrews, and Ziebarth (2019), linking patent records with historical US census records), and Aghion, Akcigit, Hyytinen, and Toivanen (2018) for modern Finnish data (linking patent records to Finnish registry data). Perhaps most strikingly, these three papers uncover a very similar relationship between parental resources and the probability of an individual becoming an inventor. As shown in Figure 16 (using data from Akcigit, Grigsby, and Nicholas (2017b), Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019c), and Aghion, Akcigit, Hyytinen, and Toivanen (2018)), the probability of patenting for an individual whose father is at the very top of the income distribution is about ten times larger than the corresponding probability for an individual with a father at the bottom end of the income distribution. Aghion, Akcigit, Hyytinen, and Toivanen (2018) discuss how the similarity between Finland and the US is all the more remarkable given that Finland displays much lower income inequality and higher social mobility than the US, and offers free education up to and including university.

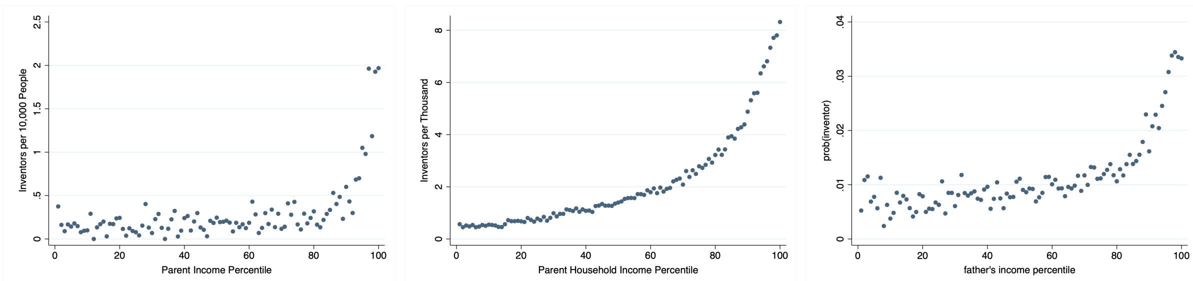

(a) Akcigit, Grigsby, and Nicholas (2017b) **(b)** Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019c) **(c)** Aghion, Akcigit, Hyytinen, and Toivanen (2018)

Figure 16: The relationship between parents' income and becoming an inventor

Notes: Panel (a) replicates Figure 21 from Akcigit, Grigsby, and Nicholas (2017b), which links individual level data on patenting (the definition of 'inventor') to fathers' percentile of wage income in the 1940 census. Panel (b) replicates Figure 1 from Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019c), which links individual level data on patent applications and grants (the definition of 'inventor') to parents' income percentile. Panel (c) replicates Figure 1 panel C from Akcigit, Grigsby, Nicholas, and Stantcheva (2018), which links patent grant data (the definition of 'inventor') to fathers' percentile of wage income for Finnish men.

Of course, descriptive statistics on differences in the likelihood of inventing across demographic groups are consistent with a number of potential explanations. Sharp tests for causal relationships in this area are challenging to construct particularly given that many demographic characteristics of interest are not amenable to quasi-experimental research approaches. Yet several papers have nonetheless managed to design thoughtful and insightful empirical tests to shed light on this set of issues. Ellison and Swanson (2010) uses data from the American Mathematics Competitions to document a striking fact: whereas top-scoring boys come from a variety of backgrounds, top-scoring girls are drawn nearly exclusively from a very small set of elite schools. This fact is consistent with the idea that almost all girls with the ability to reach high levels of math achievement are not doing so. Agarwal and Gaule (2020) compares how high school students from different countries who achieve the same scores in the International Mathematical Olympiads (IMO) differ in their subsequent likelihood of contributing to mathematics research. The authors document that – for example – a participant from a low-income country produces 34% fewer mathematics publications and receives 56% fewer mathematics citations than an equally talented participant (as measured by IMO scores) from a high-income country. Like the Ellison and Swanson result, while this fact allows for multiple interpretations, it suggests that many talented would-be mathematicians from low-income countries do not achieve their mathematical career potential. Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019a) uses a variety of empirical approaches – such as analyzing children whose families move across high-innovation and low-innovation areas – to argue in favor of a causal relationship between "exposure" to innovation during childhood and children's eventually propensity to innovate themselves. Kügler (2021) and Khan (2009) both suggest that high patent fees led to patent-filing inventors being drawn from upper socioeconomic classes.

While each of these results is suggestive that changes in public policies could increase the supply of innovators from less-advantaged backgrounds, it is hard to infer the quantitative importance of such potential changes from these descriptive analyses.[110] One paper that aims to overcome that problem is the work of Hsieh, Hurst, Jones, and Klenow (2019), which starts with the observation that over the past several decades, there has been a great deal of convergence in occupations across gender and race. For example, while in 1960 94 percent of doctors and lawyers were white men, by 2010 that share was around 60 percent – suggesting that a substantial pool of women and non-white men may not have been pursuing their comparative advantage in 1960. The authors leverage a Roy (1951)-inspired model of occupational choice to empirically assess the aggregate consequences of this (mis)allocation of talent, and estimate that between 20-40% of the increase in output per person between 1960 and 2010 can be explained by an improved allocation of talent. While of course this type of estimate relies on a number of assumptions, this type of work is incredibly valuable for gauging the potential importance of this channel.

6 Conclusion

It has now been over sixty years since the 1962 NBER Rate and Direction meeting, and seventy years since the 1951 Quantitative Description of Technological Change conference. The relative importance of some topics has shifted since those pioneering events: medical innovation has supplanted Cold War weapons development as a key topic, for one (Nelson, 1962). Nonetheless, the problems identified by those early economists of innovation remain at the core of the field. Fixed costs borne only by the first inventor, knowledge spillovers, and unhedged uncertainty combine to make innovation, and hence growth, suboptimal in laissez faire. Non-market policies helping science function more efficiently, or market policies lowering the cost or increasing the return to R&D can push innovation closer to the first best. Once inventions are discovered, frictions in their diffusion must still be overcome. And policymakers must pay attention to how these innovations affect inequality, or worse, how inequality leads great inventions to go missing.

Continued progress on understanding the effects of innovation policy relies on developments in three areas. First, we still require better measurement of the "paper trail" of knowledge, of who knew what and when, and of who was using this knowledge in their processes or new products. This better measurement will come from novel techniques for analyzing and comparing large text corpora, expanded large-scale surveys of inventors and firms, linked datasets connecting R&D with pure science with patents with new products, and clever applications of datasets other than patents and academic articles. Second, advances in statistical identification and productivity measurement are needed to answer questions including some of the very most basic in all of innovation. For example, which fields of science have the highest social return? The answer to this question requires solving both a series of very tricky measurement problems and finding suitable empirical variation in funding across fields. Indeed, it is particularly challenging to measure the return to innovation in the financial sector, business processes, construction and retail, despite the increasing importance of these sectors relative to the product market.[111] Third, theoretical advances are still required to help guide these empirical investigations. Of particular importance are open questions on the microfoundations of diffusion frictions, competition policy in innovative industries, the division of rent between inventors of complements,

[110] Two recent papers have investigated the related question of whether the gender composition of innovators matters for the direction of innovation: Koning, Samila, and Ferguson (2020) and Einiö, Feng, and Jaravel (2019).

[111] See, e.g., Holmes (2011) on the importance of store density to the gains from Walmart's business process innovations. In construction, Mowery and Rosenberg (1991) gives a detailed qualitative description of how testing and standards permitted more efficient use of lumber by reducing the need for "overengineering."

and methods other than patents for credibly rewarding breakthroughs in areas unanticipated by the planner.

Innovation has long been an interdisciplinary field, drawing from economics, sociology, history, law, and management, among others. The empirical and technical breakthroughs in the economics of education, health, public finance, and industrial organization offer hope that some of the thorniest problems in the economics of innovation are similarly amenable. It is hard to overstate the importance of progress on these problems. Innovation is the tool which allows society to escape the bonds of scarcity. Without being constrained by scarce resources, our most fundamental challenges - poverty, climate change, inequality - are solvable.

References

Abrams, David S. 2009. "Did TRIPS Spur Innovation? An Empirical Analysis of Patent Duration and Incentives to Innovate." *University of Pennsylvania Law Review* 157: 1613–1647.

Abrams, David S., Ufuk Akcigit, and Jillian Grennan. 2019. "Patent Value and Citations: Creative Destruction or Strategic Disruption?" Institute for Law & Econ Working Paper 13-23. University of Pennsylvania.

Acemoglu, Daron. 2002. "Directed Technical Change." *Review of Economic Studies* 69 (4): 781–809.

— 2012. "Diversity and Technological Progress." In: *The Rate and Direction of Inventive Activity Revisited*. Chicago, IL: University of Chicago Press: 319–356.

Acemoglu, Daron and Ufuk Akcigit. 2012. "Intellectual Property Rights Policy, Competition and Innovation." *Journal of the European Economic Association* 10 (1): 1–42.

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr. 2016. "Transition to Clean Technology." *Journal of Political Economy* 124 (1): 52–104.

Acemoglu, Daron and Joshua Linn. 2004. "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry." *Quarterly Journal of Economics* 119 (3): 1049–1090.

Adam, David. 2019. "Science Funders Gamble on Grant Lotteries." *Nature* 575 (7784): 574–575.

Adams, James D. and J. Roger Clemmons. 2013. "How Rapidly Does Science Leak Out? A Study of the Diffusion of Fundamental Ideas." *Journal of Human Capital* 7 (3): 191–229.

Adhvaryu, Achyuta. 2014. "Learning, Misallocation, and Technology Adoption: Evidence from New Malaria Therapy in Tanzania." *Review of Economic Studies* 81 (4): 1331–1365.

Agarwal, Ruchir, Ina Ganguli, Patrick Gaulé, and Geoff Smith. 2021. "Why U.S. Immigration Barriers Matter for the Global Advancement of Science." IZA Discussion Paper 14016. IZA Institute of Labor Economics.

Agarwal, Ruchir and Patrick Gaule. 2020. "Invisible Geniuses: Could the Knowledge Frontier Advance Faster?" *American Economic Review: Insights* 2 (4): 409–424.

Agha, Leila and David Molitor. 2018. "The Local Influence of Pioneer Investigators on Technology Adoption: Evidence from New Cancer Drugs." *Review of Economics and Statistics* 100 (1): 29–44.

Aghion, Philippe, Ufuk Akcigit, Antonin Bergeaud, Richard Blundell, and David Hemous. 2019. "Innovation and Top Income Inequality." *Review of Economic Studies* 86 (1): 1–45.

Aghion, Philippe, Ufuk Akcigit, Ari Hyytinen, and Otto Toivanen. 2018. "The Social Origins and IQ of Inventors." NBER Working Paper 24110. National Bureau of Economic Research.

Aghion, Philippe, Nicholas Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. "Competition and Innovation: An Inverted-U Relationship." *Quarterly Journal of Economics* 120 (2): 701–728.

Aghion, Philippe, Mathias Dewatripont, and Jeremy C. Stein. 2008. "Academic Freedom, Private-Sector Focus, and the Process of Innovation." *RAND Journal of Economics* 39 (3): 617–635.

Aghion, Philippe and Rachel Griffith. 2008. *Competition and Growth: Reconciling Theory and Evidence*. Cambridge, MA: MIT Press.

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

Aghion, Philippe and Jean Tirole. 1994. "The Management of Innovation." *Quarterly Journal of Economics* 109 (4): 1185–1209.

Aghion, Phillipe, Chad Jones, and Benjamin F. Jones. 2019. "Artificial Intelligence and Economic Growth." In: *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press: 237–282.

Agrawal, Ajay, Alberto Galasso, and Alexander Oettl. 2017. "Roads and Innovation." *Review of Economics and Statistics* 99 (3): 417–434.

Agrawal, Ajay, Avi Goldfarb, and Florenta Teodoridis. 2016. "Understanding the Changing Structure of Scientific Inquiry." *American Economic Journal: Applied Economics* 8 (1): 100–128.

Agrawal, Ajay, Carlos Rosell, and Timothy Simcoe. 2020. "Tax Credits and Small Firm R&D Spending." *American Economic Journal: Economic Policy* 12 (2): 1–21.

Ahmadpoor, Mohammad and Benjamin F. Jones. 2017. "The Dual Frontier: Patented Inventions and Prior Scientific Advance." *Science* 357 (6351): 583–587.

Akbarpour, Mohammad, Suraj Malladi, and Amin Saberi. 2020. "Just a Few Seeds More: Value of Network Information for Diffusion." SSRN Scholarly Paper ID 3062830. Social Science Research Network.

Akcigit, Ufuk, Salomé Baslandze, and Stefanie Stantcheva. 2016. "Taxation and the International Mobility of Inventors." *American Economic Review* 106 (10): 2930–2981.

Akcigit, Ufuk, John Grigsby, and Tom Nicholas. 2017a. "Immigration and the Rise of American Ingenuity." *American Economic Review* 107 (5): 327–331.

— 2017b. "The Rise of American Ingenuity: Innovation and Inventors of the Golden Age." NBER Working Paper 23047. National Bureau of Economic Research.

Akcigit, Ufuk, John Grigsby, Tom Nicholas, and Stefanie Stantcheva. 2018. "Taxation and Innovation in the 20th Century." NBER Working Paper 24982. National Bureau of Economic Research.

Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde. Forthcoming. "Back to Basics: Basic Research Spillovers, Innovation Policy and Growth." *Review of Economic Studies*.

Akcigit, Ufuk, Douglas Hanley, and Stefanie Stantcheva. 2019. "Optimal Taxation and R&D Policies." NBER Working Paper 22908. Cambridge, MA: National Bureau of Economic Research.

Alcácer, Juan, Michelle Gittelman, and Bhaven N. Sampat. 2009. "Applicant and Examiner Citations in U.S. Patents: An Overview and Analysis." *Research Policy* 38 (2): 415–427.

Allen, Robert C. 1983. “Collective Invention.” *Journal of Economic Behavior & Organization* 4 (1): 1–24.

— 2009. “The Industrial Revolution in Miniature: The Spinning Jenny in Britain, France, and India.” *Journal of Economic History* 69 (4): 901–927.

Allen, Treb, L. Kamran Bilir, Zhang Chen, and Christopher Tonetti. 2019. “Knowledge Diffusion Through Networks.” SSRN Scholarly Paper 3509835. Social Science Research Network.

Allison, John R., Mark A. Lemley, Kimberly A. Moore, and R. Derek Trunkey. 2004. “Valuable Patents.” *Georgetown Law Journal* 92 (3): 435–480.

Allison, John R., Mark A. Lemley, and Joshua Walker. 2009. “Extreme Value or Trolls on Top? The Characteristics of the Most-Litigated Patents.” *University of Pennsylvania Law Review* 158 (1): 1–37.

Alstadsæter, Annette, Salvador Barrios, Gaetan Nicodeme, Agnieszka Maria Skonieczna, and Antonio Vezzani. 2018. “Patent Boxes Design, Patents Location, and Local R&D.” *Economic Policy* 33 (93): 131–177.

Alston, Julian M. and Philip G. Pardey. 2014. “Agriculture in the Global Economy.” *Journal of Economic Perspectives* 28 (1): 121–146.

Andrews, Michael J. 2021. “Historical Patent Data: A Practitioner’s Guide.” *Journal of Economics & Management Strategy* 30 (2): 368–397.

Anton, James J. and Dennis A. Yao. 1994. “Expropriation and Inventions: Appropriable Rents in the Absence of Property Rights.” *American Economic Review* 84 (1): 190–209.

— 2004. “Little Patents and Big Secrets: Managing Intellectual Property.” *RAND Journal of Economics* 35 (1): 1–22.

Appelt, Silvia, Matej Bajgar, Chiara Criscuolo, and Fernando Galindo-Rueda. 2016. “R&D Tax Incentives: Evidence on Design, Incidence and Impacts.” OECD Science, Technology, and Industry Policy Papers No. 32. OECD.

Argente, David, Salomé Baslandze, Douglas Hanley, and Sara Moreira. 2020. “Patents to Products: Product Innovation and Firm Dynamics.” FRB Atlanta Working Paper 2020-4. Federal Reserve Bank of Atlanta.

Arora, Ashish, Sharon Belenzon, and Honggi Lee. 2018. “Reversed Citations and the Localization of Knowledge Spillovers.” *Journal of Economic Geography* 18 (3): 495–521.

Arora, Ashish, Sharon Belenzon, and Andrea Patacconi. 2018. “The Decline of Science in Corporate R&D.” *Strategic Management Journal* 39 (1): 3–32.

Arora, Ashish, Sharon Belenzon, Andrea Patacconi, and Jungkyu Suh. 2019. “The Changing Structure of American Innovation: Some Cautionary Remarks for Economic Growth.” *Innovation Policy and the Economy* 20: 39–93.

Arora, Ashish, Sharon Belenzon, and Lia Sheer. 2021. “Knowledge Spillovers and Corporate Investment in Scientific Research.” *American Economic Review* 111 (3): 871–898.

Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella. 2001. *Markets for Technology: The Economics of Innovation and Corporate Strategy*. Cambridge, MA: MIT Press.

Arora, Ashish and Alfonso Gambardella. 2010. "The Market for Technology." In: *Handbook of the Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 1. North-Holland: 641–678.

Arque-Castells, Pere and Daniel F. Spulber. 2019. "Measuring the Private and Social Returns to R&D: Unintended Spillovers versus Technology Markets." Law & Econ Research Paper 18-18. Northwestern University.

Arrow, Kenneth J. 1962a. "Economic Welfare and the Allocation of Resources for Invention." In: *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press: 609–626.

— 1962b. "The Economic Implications of Learning by Doing." *Review of Economic Studies* 29 (3): 155–173.

— 1969. "Classificatory Notes on the Production and Transmission of Technological Knowledge." *American Economic Review* 59 (2): 29–35.

Arrow, Kenneth J., L. Kamran Bilir, and Alan Sorensen. 2020. "The Impact of Information Technology on the Diffusion of New Pharmaceuticals." *American Economic Journal: Applied Economics* 12 (3): 1–39.

Arthur, W. Brian. 1989. "Competing Technologies, Increasing Returns, and Lock-In by Historical Events." *Economic Journal* 99 (394): 116–131.

Asher, Sam and Paul Novosad. 2020. "Rural Roads and Local Economic Development." *American Economic Review* 110 (3): 797–823.

Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K. Khandelwal, and Eric Verhoogen. 2017. "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan." *Quarterly Journal of Economics* 132 (3): 1101–1164.

Audretsch, David B. and Maryann P. Feldman. 1996. "R&D Spillovers and the Geography of Innovation and Production." *American Economic Review* 86 (3): 630–640.

Autor, David, David Dorn, Gordon H. Hanson, Gary Pisano, and Pian Shu. 2020. "Foreign Competition and Domestic Innovation: Evidence from US Patents." *American Economic Review: Insights* 2 (3): 357–374.

Azoulay, Pierre, Christian Fons-Rosen, and Joshua S. Graff Zivin. 2019. "Does Science Advance One Funeral at a Time?" *American Economic Review* 109 (8): 2889–2920.

Azoulay, Pierre, Joshua S. Graff Zivin, Danielle Li, and Bhaven N. Sampat. 2019. "Public R&D Investments and Private-Sector Patenting: Evidence from NIH Funding Rules." *Review of Economic Studies* 86 (1): 117–152.

Azoulay, Pierre, Joshua S. Graff Zivin, and Gustavo Manso. 2011. "Incentives and Creativity: Evidence from the Academic Life Sciences." *RAND Journal of Economics* 42 (3): 527–554.

Azoulay, Pierre, Benjamin F. Jones, J. Daniel Kim, and Javier Miranda. Forthcoming. "Immigration and Entrepreneurship in the United States." *American Economic Review: Insights.*

Azoulay, Pierre and Danielle Li. Forthcoming. "Scientific Grant Funding." *Innovation and Public Policy.*

Baker, Jonathan. 2007. "Beyond Schumpeter vs. Arrow: How Antitrust Fosters Innovation." *Antitrust Law Journal* 74 (3): 29.

Bandiera, Oriana and Imran Rasul. 2006. "Social Networks and Technology Adoption in Northern Mozambique." *Economic Journal* 116 (514): 869–902.

Barany, Michael J. 2019. "Rockefeller Bureaucracy and Circumknowing Science in the Mid-Twentieth Century." *International Journal for History, Culture and Modernity* 7 (1): 779–796.

Barzel, Yoram. 1968. "Optimal Timing of Innovations." *Review of Economics and Statistics* 50 (3): 348–355.

Basker, Emek and Timothy Simcoe. 2021. "Upstream, Downstream: Diffusion and Impacts of the Universal Product Code." *Journal of Political Economy* 129 (4): 1252–1286.

Bell, Alexander, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. 2019a. "Do Tax Cuts Produce More Einsteins? The Impacts of Financial Incentives vs. Exposure to Innovation on the Supply of Inventors." NBER Working Paper 25493. National Bureau of Economic Research.

— 2019b. "Joseph Schumpeter Lecture, EEA Annual Congress 2017: Do Tax Cuts Produce More Einsteins? The Impacts of Financial Incentives Versus Exposure to Innovation on the Supply of Inventors." *Journal of the European Economic Association* 17 (3): 651–677.

— 2019c. "Who Becomes an Inventor in America? The Importance of Exposure to Innovation." *Quarterly Journal of Economics* 134 (2): 647–713.

Benhabib, Jess, Jesse Perla, and Christopher Tonetti. 2014. "Catch-up and Fall-Back through Innovation and Imitation." *Journal of Economic Growth* 19 (1): 1–35.

Bentham, Jeremy. 1825. *The Rationale of Reward.* London: John and H.L. Hunt.

Berkes, Enrico. 2016. "Comprehensive Universe of U.S. Patents (CUSP): Data and Facts." Working Paper.

Berkes, Enrico and Ruben Gaetani. 2019. "Income Segregation and Rise of the Knowledge Economy." Rotman School of Management Working Paper No. 3423136. Rotman School of Management.

Bernstein, Jeffrey I. and M. Ishaq Nadiri. 1989. "Research and Development and Intra-Industry Spillovers: An Empirical Application of Dynamic Duality." *Review of Economic Studies* 56 (2): 249–267.

Bernstein, Shai, Rebecca Diamond, Timothy James McQuade, and Beatriz Pousada. 2019. "The Contribution of High-Skilled Immigrants to Innovation in the United States." Working Paper Working Paper No. 3748. Stanford Graduate School of Business.

Bessen, James E. 2008. “The Value of U.S. Patents by Owner and Patent Characteristics.” *Research Policy* 37 (5): 932–945.

Bessen, James E. and Eric Maskin. 2009. “Sequential Innovation, Patents, and Imitation.” *RAND Journal of Economics* 40 (4): 611–635.

Bessen, James E. and Michael J. Meurer. 2005. “Lessons for Patent Policy from Empirical Research on Patent Litigation.” *Lewis and Clark Law Review* 9 (1): 1–29.

Bhattacharya, Vivek. Forthcoming. “An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program.” *Econometrica*.

Biasi, Barbara and Petra Moser. Forthcoming. “Effects of Copyrights on Science: Evidence from the WWII Book Republication Program.” *American Economic Journal: Microeconomics*.

Bikard, Michaël. 2020. “Idea Twins: Simultaneous Discoveries as a Research Tool.” *Strategic Management Journal* 41 (8): 1528–1543.

Bilir, L. Kamran and Eduardo Morales. 2020. “Innovation in the Global Firm.” *Journal of Political Economy* 128 (24): 1566–1625.

Birkhaeuser, Dean, Robert E. Evenson, and Gershon Feder. 1991. “The Economic Impact of Agricultural Extension: A Review.” *Economic Development and Cultural Change* 39 (3): 607–650.

Block, Fred and Matthew R. Keller. 2009. “Where Do Innovations Come from? Transformations in the US Economy, 1970–2006.” *Socio-Economic Review* 7 (3): 459–483.

Bloom, Gordon F. 1946. “A Note on Hicks’s Theory of Invention.” *American Economic Review* 36 (1): 83–96.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen. 2019. “What Drives Differences in Management Practices?” *American Economic Review* 109 (5): 1648–1683.

Bloom, Nicholas, Mirko Draca, and John Van Reenen. 2016. “Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity.” *Review of Economic Studies* 83 (1): 87–117.

Bloom, Nicholas, Rachel Griffith, and John Van Reenen. 2002. “Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979-1997.” *Journal of Public Economics* 85 (1): 1–31.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2020. “Are Ideas Getting Harder to Find?” *American Economic Review* 110 (4): 1104–1144.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. “Identifying Technology Spillovers and Product Market Rivalry.” *Econometrica* 81 (4): 1347–1393.

Bloom, Nicholas and John Van Reenen. 2007. “Measuring and Explaining Management Practices Across Firms and Countries.” *Quarterly Journal of Economics* 122 (4): 1351–1408.

Bloom, Nicholas, John Van Reenen, and Heidi L. Williams. 2019. “A Toolkit of Policies to Promote Innovation.” *Journal of Economic Perspectives* 33 (3): 163–184.

Blume-Kohout, Margaret E. 2012. "Does Targeted, Disease-Specific Public Research Funding Influence Pharmaceutical Innovation?" *Journal of Policy Analysis and Management* 31 (3): 641–660.

Blume-Kohout, Margaret E. and Neeraj Sood. 2013. "Market Size and Innovation: Effects of Medicare Part D on Pharmaceutical Research and Development." *Journal of Public Economics* 97 (1): 327–336.

Bobtcheff, Catherine, Jérôme Bolte, and Thomas Mariotti. 2016. "Researcher's Dilemma." *Review of Economic Studies* 84 (3): 969–1014.

Boldrin, Michele and David K. Levine. 2013. "The Case Against Patents." *Journal of Economic Perspectives* 27 (1): 3–22.

Bonatti, Alessandro and Johannes Hörner. 2011. "Collaborating." *American Economic Review* 101 (2): 632–63.

Borjas, George J. and Kirk B. Doran. 2012. "The Collapse of the Soviet Union and the Productivity of American Mathematicians." *Quarterly Journal of Economics* 127 (3): 1143–1203.

Branstetter, Lee G., Britta Glennon, and J. Bradford Jensen. 2019. "The Rise of Global Innovation by US Multinationals Poses Risks and Opportunities." Policy Brief 19-9. Peterson Institute for International Economics.

Brendel, Jan and Sascha Schweitzer. 2019. "The Burden of Knowledge in Mathematics." *Open Economics* 2 (1): 139–149.

Bresnahan, Timothy F. 2010. "General Purpose Technologies." In: *Handbook of the Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 2. North-Holland: 761–791.

Bresnahan, Timothy F. and Manuel Trajtenberg. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* 65 (1): 83–108.

Bresnahan, Timothy F, Shane Greenstein, and Rebecca M. Henderson. 2010. "Schumpeterian Competition and Diseconomies of Scope: Illustrations from the Histories of Microsoft and IBM." In: *The Rate and Direction of Inventive Activity Revisited*. Ed. by Josh Lerner and Scott Stern. Chicago, IL: University of Chicago Press: 203–271.

Bronzini, Raffaello and Eleonora Iachini. 2014. "Are Incentives for R&D Effective? Evidence from a Regression Discontinuity Approach." *American Economic Journal: Economic Policy* 6 (4): 100–134.

Brooks, Wyatt and Kevin Donovan. 2020. "Eliminating Uncertainty in Market Access: The Impact of New Bridges in Rural Nicaragua." *Econometrica* 88 (5): 1965–1997.

Brown, Murray and Alfred H. Conrad. 1967. "The Influence of Research and Education on CES Production Relations." In: *The Theory and Empirical Analysis of Production*. New York, NY: National Bureau of Economic Research, Inc: 341–394.

Brown, William H. 1957. "Innovation in the Machine Tool Industry." *Quarterly Journal of Economics* 71 (3): 406–425.

Brozen, Yale. 1951. "Invention, Innovation, and Imitation." *American Economic Review* 41 (2): 239–257.

Bryan, Kevin A. and Erik Hovenkamp. 2020. "Antitrust Limits on Startup Acquisitions." *Review of Industrial Organization* 56 (4): 615–636.

Bryan, Kevin A. and Jorge Lemus. 2017. "The Direction of Innovation." *Journal of Economic Theory* 172: 247–272.

Bryan, Kevin A. and Yasin Ozcan. Forthcoming. "The Impact of Open Access Mandates on Invention." *Review of Economics and Statistics*.

Bryan, Kevin A., Yasin Ozcan, and Bhaven N Sampat. 2020. "In-Text Patent Citations: A User's Guide." *Research Policy* 49 (4): Article 103946.

Budish, Eric, Benjamin N. Roin, and Heidi L. Williams. 2015. "Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials." *American Economic Review* 105 (7): 2044–2085.

— 2016. "Patents and Research Investments: Assessing the Empirical Evidence." *American Economic Review* 106 (5): 183–187.

Buera, Francisco J. and Ezra Oberfield. 2020. "The Global Diffusion of Ideas." *Econometrica* 88 (1): 83–114.

Burchardi, Konrad B., Thomas Chaney, Tarek A Hassan, Lisa Tarquinio, and Stephen J. Terry. 2020. "Immigration, Innovation, and Growth." NBER Working Paper 27075. National Bureau of Economic Research.

Burk, Dan L. and Mark A. Lemley. 2003. "Policy Levers in Patent Law." *Virginia Law Review* 89 (7): 1575–1696.

Bush, Vannevar. 1945. "Science, the Endless Frontier: A Report to the President on a Program for Postwar Scientific Research." Tech. rep. Washington: Office of Scientific Research and Development.

Cabral, Luís M. B. 2003. "R&D Competition When Firms Choose Variance." *Journal of Economics & Management Strategy* 12 (1): 139–150.

— 2018. "Standing on the Shoulders of Dwarfs: Dominant Firms and Innovation Incentives." CEPR Discussion Paper No. DP13115. Centre for Economic Policy Research: 26.

Callander, Steven and Niko Matouschek. Forthcoming. "The Novelty of Innovation: Competition, Disruption, and Antitrust Policy." *Management Science*.

Carley, Michael, Deepak Hegde, and Alan C. Marco. 2015. "What Is the Probability of Receiving a U.S. Patent?" *Yale Journal of Law and Technology* 17 (1): 203–223.

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian. 2019. "Are Markups Too High? Competition, Strategic Innovation, and Industry Dynamics." SSRN Scholarly Paper ID 3459775. Rochester, NY: Social Science Research Network.

Chakravarthy, Ranjana, Kristina Cotter, Joseph DiMasi, Christopher-Paul Milne, and Nils Wendel. 2016. "Public- and Private-Sector Contributions to the Research and Development of the Most Transformational Drugs in the Past 25 Years: From Theory to Therapy." *Therapeutic Innovation & Regulatory Science* 50 (6): 759–768.

Chamberlin, Edward Hastings. 1962. *The Theory of Monopolistic Competition.* Cambridge, MA: Harvard University Press.

Chapin, F. Stuart. 1928. *Cultural Change.* New York, NY: The Century Co.

Chatterji, Aaron K. and Kira R. Fabrizio. 2014. "Using Users: When Does External Knowledge Enhance Corporate Product Innovation?" *Strategic Management Journal* 35 (10): 1427–1445.

Chawla, Dalmeet Singh. 2020. "They Wanted Research Funding, So They Entered the Lottery." *The New York Times.*

Che, Yeon-Koo and Seung-Weon Yoo. 2001. "Optimal Incentives for Teams." *American Economic Review* 91 (3): 525–541.

Chen, Yongmin and Thitima Puttitanun. 2005. "Intellectual Property Rights and Innovation in Developing Countries." *Journal of Development Economics* 78 (2): 474–493.

Chen, Zhao, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Yi Xu. 2021. "Notching R&D Investment with Corporate Income Tax Cuts in China." *American Economic Review* 111 (7): 2065–2100.

Chien, Colleen V. 2011. "Articles Predicting Patent Litigation." *Texas Law Review* 90 (283): 284–329.

Clark, John Bates. 1915. *Essentials of Economic Theory.* New York, NY: The Macmillan Company.

Cleary, Ekaterina Galkina, Jennifer M. Beierlein, Navleen Surjit Khanuja, Laura M. McNamee, and Fred D. Ledley. 2018. "Contribution of NIH Funding to New Drug Approvals 2010–2016." *Proceedings of the National Academy of Sciences* 115 (10): 2329–2334.

Cockburn, Iain M., Samuel Kortum, and Scott Stern. 2003. "Are All Patent Examiners Equal? The Impact of Examiner Characteristics." In: *Patents in the Knowledge-Based Economy.* Ed. by Wesley Cohen and Cohen Merrill. National Academies Press.

Cockburn, Iain M., Jean O. Lanjouw, and Mark Schankerman. 2016. "Patents and the Global Diffusion of New Drugs." *American Economic Review* 106 (1): 136–164.

Coe, David T., Elhanan Helpman, and Alexander W. Hoffmaister. 2009. "International R&D Spillovers and Institutions." *European Economic Review* 53 (7): 723–741.

Cohen, Wesley M. 2010. "Fifty Years of Empirical Studies of Innovative Activity and Performance." In: *Handbook of the Economics of Innovation.* Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 1. North-Holland: 129–213.

Cohen, Wesley M. and Steven Klepper. 1996. "A Reprise of Size and R & D." *Economic Journal* 106 (437): 925–951.

Cohen, Wesley M. and Richard C. Levin. 1989. "Chapter 18 Empirical Studies of Innovation and Market Structure." *Handbook of Industrial Organization* 2: 1059–1107.

Cohen, Wesley M. and Daniel A. Levinthal. 1989. "Innovation and Learning: The Two Faces of R & D." *Economic Journal* 99 (397): 569–596.

Cohen, Wesley M., Richard R. Nelson, and John P. Walsh. 2000. "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)." NBER Working Paper 7552. National Bureau of Economic Research.

Cohen, Wesley M. and John P. Walsh. 2007. "Real Impediments to Academic Biomedical Research." *Innovation Policy and the Economy* 8 (1): 1–30.

Coleman, James, Elihu Katz, and Herbert Menzel. 1957. "The Diffusion of an Innovation Among Physicians." *Sociometry* 20 (4): 253–270.

Collison, Patrick, Tyler Cowen, and Patrick Hsu. 2021. *What We Learned Doing Fast Grants.* https://future.a16z.com/what-we-learned-doing-fast-grants/.

Comin, Diego, William Easterly, and Erick Gong. 2010. "Was the Wealth of Nations Determined in 1000 BC?" *American Economic Journal: Macroeconomics* 2 (3): 65–97.

Comin, Diego and Bart Hobijn. 2010. "An Exploration of Technology Diffusion." *American Economic Review* 100 (5): 2031–2059.

Comin, Diego and Martí Mestieri. 2014. "Technology Diffusion: Measurement, Causes, and Consequences." In: *Handbook of Economic Growth.* Ed. by Philippe Aghion and Steven N. Durlauf. Vol. 2. Elsevier: 565–622.

— 2018. "If Technology Has Arrived Everywhere, Why Has Income Diverged?" *American Economic Journal: Macroeconomics* 10 (3): 137–178.

Comino, Stefano, Alberto Galasso, and Clara Graziano. 2020. "Market Power and Patent Strategies: Evidence from Renaissance Venice." *Journal of Industrial Economics* 68 (2): 226–269.

Conley, Timothy G. and Christopher R. Udry. 2010. "Learning about a New Technology: Pineapple in Ghana." *American Economic Review* 100 (1): 35–69.

Cook, Lisa D. 2014. "Violence and Economic Activity: Evidence from African American Patents, 1870-1940." *Journal of Economic Growth* 19 (2): 221–257.

Cook, Lisa D. and Chaleampong Kongcharoen. 2010. "The Idea Gap in Pink and Black." NBER Working Paper 16331. National Bureau of Economic Research.

Costa, Dora L. and Naomi R. Lamoreaux, eds. 2011. *Understanding Long-Run Economic Growth.* National Bureau of Economic Research Conference Report. Chicago, IL: University of Chicago Press.

Costinot, Arnaud, Dave Donaldson, Margaret K. Kyle, and Heidi L. Williams. 2019. "The More We Die, The More We Sell? A Simple Test of the Home-Market Effect." *Quarterly Journal of Economics* 134 (2): 843–894.

Cowan, Robin. 1990. "Nuclear Power Reactors: A Study in Technological Lock-In." *Journal of Economic History* 50 (3): 541–567.

Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer. 2019. "Competition Policy for the Digital Era." Technical Report. Publications Office of the European Union.

Cremers, Katrin. 2004. "Determinants of Patent Litigation in Germany." Discussion Paper 04-072. ZEW - Centre for European Economic Research.

Cunningham, Colleen, Florian Ederer, and Song Ma. 2021. "Killer Acquisitions." *Journal of Political Economy* 129 (3): 649–702.

Curello, Gregorio. 2021. "Incentives for Collective Innovation." Working Paper. Institute for Microeconomics, University of Bonn.

Cutler, David M., Ellen Meara, and Seth Richards-Shubik. 2012. "Induced Innovation and Social Inequality: Evidence from Infant Medical Care." *Journal of Human Resources* 47 (2): 456–492.

Dasgupta, Partha. 1989. "The Economics of Parallel Research." In: *The Economics of Missing Markets, Information and Games*. Ed. by Frank Hahn. Oxford: Clarendon Press: 129–148.

Dasgupta, Partha and Paul A. David. 1987. "Information Disclosure and the Economics of Science and Technology." In: *Arrow and the Ascent of Modern Economic Theory*. Ed. by George R. Feiwel. London: Palgrave Macmillan UK: 519–542.

— 1994. "Toward a New Economics of Science." *Research Policy* 23 (5): 487–521.

Dasgupta, Partha and Eric Maskin. 1987. "The Simple Economics of Research Portfolios." *Economic Journal* 97 (387): 581–595.

David, Paul A. 1985. "Clio and the Economics of QWERTY." *American Economic Review* 75 (2): 332–337.

— 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80 (2): 355–361.

David, Paul A., Bronwyn H. Hall, and Andrew A. Toole. 2000. "Is Public R&D a Complement or Substitute for Private R&D? A Review of the Econometric Evidence." *Research Policy* 29 (4): 497–529.

de Rassenfosse, Gaétan and Adam B. Jaffe. 2018. "Are Patent Fees Effective at Weeding Out Low-Quality Patents?" *Journal of Economics & Management Strategy* 27 (1): 134–148.

de Rassenfosse, Gaétan, Adam B. Jaffe, and Emilio Raiteri. 2019. "The Procurement of Innovation by the U.S. Government." *PLOS ONE* 14 (8): Article 0218927.

de Rassenfosse, Gaétan, Adam B Jaffe, and Elizabeth Webster. 2016. "Low-Quality Patents in the Eye of the Beholder: Evidence from Multiple Examiners." NBER Working Paper 22244.

de Rassenfosse, Gaétan and Ling Zhou. 2021. "A Test of the Monopoly Pricing Hypothesis of Patents." SSRN Scholarly Paper ID 3756359. Rochester, NY: Social Science Research Network.

de Souza, Gustavo. 2020. "The Labor Market Consequences of Appropriate Technology." Working Paper.

Deardorff, Alan V. 1992. "Welfare Effects of Global Patent Protection." *Economica* 59 (233): 35–51.

Dechezleprêtre, Antoine, Elias Einiö, Ralf Martin, Kieu-Trang Nguyen, and John Van Reenen. 2016. "Do Tax Incentives for Research Increase Firm Innovation? An RD Design for R&D." NBER Working Paper 22405. National Bureau of Economic Research.

Denicolò, Vincenzo. 2000. "Two-Stage Patent Races and Patent Policy." *RAND Journal of Economics* 31 (3): 488–501.

Ding, Waverly W., Fiona Murray, and Toby E. Stuart. 2006. "Gender Differences in Patenting in the Academic Life Sciences." *Science* 313 (5787): 665–667.

Dittmar, Jeremiah E. 2011. "Information Technology and Economic Change: The Impact of The Printing Press." *Quarterly Journal of Economics* 126 (3): 1133–1172.

Diwan, Ishac and Dani Rodrik. 1991. "Patents, Appropriate Technology, and North-South Trade." *Journal of International Economics* 30 (1-2): 27–47.

Doran, Kirk, Alexander Gelber, and Adam Isen. 2014. "The Effects of High-Skilled Immigration Policy on Firms: Evidence from H-1B Visa Lotteries." NBER Working Paper 20668. National Bureau of Economic Research.

Doran, Kirk and Chungeun Yoon. 2020. "Immigration and Invention: Evidence from the Quota Acts." Working Paper. University of Notre Dame.

Doraszelski, Ulrich. 2001. "The Net Present Value Method versus the Option Value of Waiting: A Note on Farzin, Huisman and Kort (1998)." *Journal of Economic Dynamics and Control* 25: 1109–1115.

— 2003. "An R&D Race with Knowledge Accumulation." *RAND Journal of Economics* 34 (1): 20–42.

Dosi, Giovanni and Richard R. Nelson. 2010. "Technical Change and Industrial Dynamics as Evolutionary Processes." In: *Handbook of the Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 1. North-Holland: 51–127.

Dubois, Pierre, Olivier de Mouzon, Fiona Scott-Morton, and Paul Seabright. 2015. "Market Size and Pharmaceutical Innovation." *RAND Journal of Economics* 46 (4): 844–871.

Dudley, Leonard. 2017. "Language Standardization and the Industrial Revolution." *Oxford Economic Papers* 69 (4): 1138–1161.

Duesenberry, James. 1956. "Innovation and Growth." *American Economic Review* 46 (2): 134–141.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2008. "How High Are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya." *American Economic Review* 98 (2): 482–488.

Duggan, Mark, Craig Garthwaite, and Aparajita Goyal. 2016. "The Market Impacts of Pharmaceutical Product Patents in Developing Countries: Evidence from India." *American Economic Review* 106 (1): 99–135.

Durvasula, Maya, Lisa Larrimore Ouellette, and Heidi Williams. 2021. "Private and Public Investments in Biomedical Research." *AEA Papers and Proceedings* 111: 341–345.

Eaton, Jonathan and Samuel Kortum. 1999. "International Technology Diffusion: Theory and Measurement." *International Economic Review* 40 (3): 537–570.

Einiö, Elias. 2014. "R&D Subsidies and Company Performance: Evidence from Geographic Variation in Government Funding Based on the ERDF Population-Density Rule." *Review of Economics and Statistics* 96 (4): 710–728.

Einiö, Elias, Josh Feng, and Xavier Jaravel. 2019. "Social Push and the Direction of Innovation." SSRN Scholarly Paper 3383703. Rochester, NY: Social Science Research Network.

Eisenberg, Rebecca S. 1996. "Public Research and Private Development: Patents and Technology Transfer in Government-Sponsored Research." *Virginia Law Review* 82 (8): 1663–1727.

Ellison, Glenn, Edward L. Glaeser, and William R. Kerr. 2010. "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns." *American Economic Review* 100 (3): 1195–1213.

Ellison, Glenn and Ashley Swanson. 2010. "The Gender Gap in Secondary School Mathematics at High Achievement Levels: Evidence from the American Mathematics Competitions." *Journal of Economic Perspectives* 24 (2): 109–128.

Epstein, Ralph C. 1926. "Industrial Invention: Heroic, or Systematic?" *Quarterly Journal of Economics* 40 (2): 232–272.

Ewens, Michael and Richard R. Townsend. 2020. "Are Early Stage Investors Biased against Women?" *Journal of Financial Economics* 135 (3): 653–677.

Farrell, Joseph and Paul Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." In: *Handbook of Industrial Organization*. Ed. by M. Armstrong and R. Porter. Vol. 3. Elsevier: 1967–2072.

Farrell, Joseph and Garth Saloner. 1986. "Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation." *American Economic Review* 76 (5): 940–955.

Farzin, Y. H., K. J. M. Huisman, and P. M. Kort. 1998. "Optimal Timing of Technology Adoption." *Journal of Economic Dynamics and Control* 22 (5): 779–799.

Feng, Josh and Xavier Jaravel. 2020. "Crafting Intellectual Property Rights: Implications for Patent Assertion Entities, Litigation, and Innovation." *American Economic Journal: Applied Economics* 12 (1): 140–181.

Fernald, John G. 2012. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." FRBSF Working Paper 2012-19. Federal Reserve Bank of San Francisco: 01–28.

Finkelstein, Amy. 2004. “Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry.” *Quarterly Journal of Economics* 119 (2): 527–564.

Foster, Andrew D. and Mark R. Rosenzweig. 1995. “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture.” *Journal of Political Economy* 103 (6): 1176–1209.

— 2010. “Microeconomics of Technology Adoption.” *Annual Review of Economics* 2 (1): 395–424.

Franco, April M. and Matthew F. Mitchell. 2008. “Covenants Not to Compete, Labor Mobility, and Industry Dynamics.” *Journal of Economics & Management Strategy* 17 (3): 581–606.

Frankel, Alex. 2021. “Selecting Applicants.” *Econometrica* 89 (2): 615–645.

Furman, Jeffrey L., Markus Nagler, and Martin Watzinger. Forthcoming. “Disclosure and Subsequent Innovation: Evidence from the Patent Depository Library Program.” *American Economic Journal: Economic Policy.*

Furman, Jeffrey L. and Scott Stern. 2011. “Climbing atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research.” *American Economic Review* 101 (5): 1933–1963.

Furman, Jeffrey L. and Florenta Teodoridis. 2020. “Automation, Research Technology, and Researchers’ Trajectories: Evidence from Computer Science and Electrical Engineering.” *Organization Science* 31 (2): 330–354.

Gaessler, Fabian, Bronwyn H. Hall, and Dietmar Harhoff. 2018. “Should There Be Lower Taxes on Patent Income?” *Research Policy* 50 (1): Article 104129.

Galasso, Alberto. 2020. “Rewards versus Intellectual Property Rights When Commitment Is Limited.” *Journal of Economic Behavior and Organization* 169: 397–411.

Galasso, Alberto and Mark Schankerman. 2015. “Patents and Cumulative Innovation: Causal Evidence from the Courts.” *Quarterly Journal of Economics* 130 (1): 317–369.

Gallini, Nancy T. 1992. “Patent Policy and Costly Imitation.” *RAND Journal of Economics* 23 (1): 52–63.

Gallini, Nancy T. and Suzanne Scotchmer. 2002. “Intellectual Property: When Is It the Best Incentive System?” In: *Innovation Policy and the Economy.* Ed. by Adam B. Jaffe, Josh Lerner, and Scott Stern. Vol. 2. Cambridge, MA: MIT Press: 51–78.

Ganguli, Ina. 2017. “Saving Soviet Science: The Impact of Grants When Government R&D Funding Disappears.” *American Economic Journal: Applied Economics* 9 (2): 165–201.

Gans, Joshua S., David H. Hsu, and Scott Stern. 2002. “When Does Start-Up Innovation Spur the Gale of Creative Destruction?” *RAND Journal of Economics* 33 (4): 571–586.

Gans, Joshua S. and Scott Stern. 2000. “Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction.” *Journal of Economics & Management Strategy* 9 (4): 485–511.

Gaulé, Patrick and Mario Piacentini. 2013. “Chinese Graduate Students and U.S. Scientific Productivity.” *Review of Economics and Statistics* 95 (2): 698–701.

Geroski, Paul A. 1990. "Innovation, Technological Opportunity, and Market Structure." *Oxford Economic Papers* 42 (3): 586–602.

Gibbons, Robert and Rebecca Henderson. 2011. "Relational Contracts and Organizational Capabilities." *Organization Science* 23 (5): 1350–1364.

Gilbert, Richard J. 2006. "Looking for Mr. Schumpeter: Where Are We in the Competition–Innovation Debate?" In: *Innovation Policy and the Economy*. Vol. 6. Cambridge, MA: MIT Press: 159–215.

Gilbert, Richard J. and David M. G. Newbery. 1982. "Preemptive Patenting and the Persistence of Monopoly." *American Economic Review* 72 (3): 514–526.

Gilfillan, S. Colum. 1935. *Inventing the Ship: A Study of the Inventions Made in Her History Between Floating Log and Rotorship*. Chicago, IL: Follett Publishing Company.

Gilson, Ronald J. 1999. "The Legal Infrastructure of High Technology Industrial Districts: Silicon Valley, Route 128, and Covenants Not to Compete." *NYU Law Review* 74 (3): 575–629.

Giorcelli, Michela. 2019. "The Long-Term Effects of Management and Technology Transfers." *American Economic Review* 109 (1): 121–152.

Giorcelli, Michela and Petra Moser. 2020. "Copyrights and Creativity: Evidence from Italian Opera in the Napoleonic Age." *Journal of Political Economy* 128 (11): 4163–4210.

Glennon, Britta. 2020. "How Do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program." NBER Working Paper 27538. National Bureau of Economic Research.

Glitz, Albrecht and Erik Meyersson. 2020. "Industrial Espionage and Productivity." *American Economic Review* 110 (4): 1055–1103.

Godin, Benoît. 2008. "In the Shadow of Schumpeter: W. Rupert Maclaurin and the Study of Technological Innovation." *Minerva* 46 (3): 343–360.

— 2012. ""Innovation Studies": The Invention of a Specialty." *Minerva* 50 (4): 397–421.

Goettler, Ronald L. and Brett R. Gordon. 2011. "Does AMD Spur Intel to Innovate More?" *Journal of Political Economy* 119 (6): 1141–1200.

Gollin, Douglas, Casper Worm Hansen, and Asger Mose Wingender. 2021. "Two Blades of Grass: The Impact of the Green Revolution." *Journal of Political Economy* 129 (8): 2334–2384.

Gompers, Paul A. and Josh Lerner. 1999. *The Venture Capital Cycle*. Cambridge, MA: MIT Press.

Gompers, Paul A. and Sophie Q. Wang. 2017. "Diversity in Innovation." NBER Working Paper 23082. National Bureau of Economic Research.

Goolsbee, Austan. 1998. "Does Government R&D Policy Mainly Benefit Scientists and Engineers?" *American Economic Review* 88 (2): 298–302.

Gordon, Robert J. 2017. *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*. Princeton, NJ: Princeton University Press.

Gort, Michael and Steven Klepper. 1982. "Time Paths in the Diffusion of Product Innovations." *Economic Journal* 92 (367): 630–653.

Graham, Stuart J. H., Alan C. Marco, and Amanda F. Myers. 2018. "Patent Transactions in the Marketplace: Lessons from the USPTO Patent Assignment Dataset." *Journal of Economics & Management Strategy* 27 (3): 343–371.

Graham, Stuart J. H., Robert P. Merges, Pam Samuelson, and Ted Sichelman. 2009. "High Technology Entrepreneurs and the Patent System: Results of the 2008 Berkeley Patent Survey." *Berkeley Technology Law Journal* 24 (4): 1255.

Green, Jerry R. and Suzanne Scotchmer. 1995. "On the Division of Profit in Sequential Innovation." *RAND Journal of Economics* 26 (1): 20–33.

Greenstein, Shane and Garey Ramey. 1998. "Market Structure, Innovation and Vertical Product Differentiation." *International Journal of Industrial Organization* 16 (3): 285–311.

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

— 1962. "Comment." In: *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press: 346–353.

— 1989. "Patents: Recent Trends and Puzzles." *Brookings Papers on Economic Activity: Microeconomics* 291–319.

— 1990. "Patent Statistics as Economic Indicators: A Survey." *Journal of Economic Literature* 28 (4): 1661–1707.

— 1992. "The Search for R&D Spillovers." *Scandinavian Journal of Economics* 94: S29–S47.

Gross, Daniel P. 2018. "Scale versus Scope in the Diffusion of New Technology: Evidence from the Farm Tractor." *RAND Journal of Economics* 49 (2): 427–452.

— 2020. "Creativity under Fire: The Effects of Competition on Creative Production." *Review of Economics and Statistics* 102 (3): 583–599.

Gross, Daniel P and Bhaven N Sampat. 2020. "Inventing the Endless Frontier: The Effects of the World War II Research Effort on Post-War Innovation." NBER Working Paper 27375. Cambridge, MA: National Bureau of Economic Research.

Grossman, Gene M. and Elhanan Helpman. 1991. "Quality Ladders in the Theory of Growth." *Review of Economic Studies* 58 (1): 43–61.

Gush, Jason, Adam B. Jaffe, Victoria Larsen, and Athene Laws. 2015. "The Effect of Public Funding on Research Output: The New Zealand Marsden Fund." NBER Working Paper 21652. National Bureau of Economic Research.

Guzman, Jorge and Scott Stern. 2020. "The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 32 US States, 1988–2014." *American Economic Journal: Economic Policy* 12 (4): 212–243.

Hagiu, Andrei and David B. Yoffie. 2013. "The New Patent Intermediaries: Platforms, Defensive Aggregators, and Super-Aggregators." *Journal of Economic Perspectives* 27 (1): 45–66.

Halac, Marina, Navin Kartik, and Qingmin Liu. 2016. "Optimal Contracts for Experimentation." *Review of Economic Studies* 83 (3): 1040–1091.

— 2017. "Contests for Experimentation." *Journal of Political Economy* 125 (5): 1523–1569.

Hall, Bronwyn H. Forthcoming. "Tax Policy for Innovation." In: *Innovation and Public Policy*. Ed. by Austan Goolsbee and Benjamin Jones. University of Chicago Press.

— 1993. "R&D Tax Policy during the 1980s: Success or Failure?" *Tax Policy and the Economy* 7: 1–35.

— 2006. "Innovation and Diffusion." In: *The Oxford Handbook of Innovation*. Ed. by Jan Fagerberg and David C. Mowery. New York, NY: Oxford University Press: 459–485.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg. 2002. "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." In: *Patents, Citations and Innovations*. Ed. by Adam B Jaffe and Manuel Trajtenberg. Cambridge, MA: MIT Press.

Hall, Bronwyn H. and Josh Lerner. 2010. "The Financing of R&D and Innovation." In: *Handbook of the Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 1. North-Holland: 609–639.

Hall, Bronwyn H. and John Van Reenen. 2000. "How Effective Are Fiscal Incentives for R&D? A Review of the Evidence." *Research Policy* 29 (4): 449–469.

Hanlon, W. Walker. 2015. "Necessity Is the Mother of Invention: Input Supplies and Directed Technical Change." *Econometrica* 83 (1): 67–100.

Hargadon, Andrew B. and Yellowlees Douglas. 2001. "When Innovations Meet Institutions: Edison and the Design of the Electric Light." *Administrative Science Quarterly* 46 (3): 476–501.

Harris, Christopher and John Vickers. 1985. "Perfect Equilibrium in a Model of a Race." *Review of Economic Studies* 52 (2): 193–209.

Hegde, Deepak and Hong Luo. 2018. "Patent Publication and the Market for Ideas." *Management Science* 64 (2): 652–672.

Hellmann, Thomas. 2007. "The Role of Patents for Bridging the Science to Market Gap." *Journal of Economic Behavior and Organization*. Academic Science and Entrepreneurship: Dual Engines of Growth 63 (4): 624–647.

Helmers, Christian and Henry G. Overman. 2017. "My Precious! The Location and Diffusion of Scientific Research: Evidence from the Synchrotron Diamond Light Source." *Economic Journal* 127 (604): 2006–2040.

Helper, Susan and Rebecca Henderson. 2014. "Management Practices, Relational Contracts, and the Decline of General Motors." *Journal of Economic Perspectives* 28 (1): 49–72.

Hemel, Daniel J. and Lisa Larrimore Ouellette. 2017. “Bayh–Dole Beyond Borders.” *Journal of Law and the Biosciences* 4 (2): 282–310.

Henderson, Rebecca. 1993. “Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Equipment Industry.” *RAND Journal of Economics* 24 (2): 248–270.

Henderson, Rebecca, Adam B. Jaffe, and Manuel Trajtenberg. 2005. “Patent Citations and the Geography of Knowledge Spillovers: A Reassessment: Comment.” *American Economic Review* 95 (1): 461–464.

Henry, Emeric, Marco Loseto, and Marco Ottaviani. 2021. “Regulation with Experimentation: Ex Ante Approval, Ex Post Withdrawal, and Liability.” CEPR Discussion Paper 13224. Centre for Economic Policy Research.

Henry, Emeric and Carlos J. Ponce. 2011. “Waiting to Imitate: On the Dynamic Pricing of Knowledge.” *Journal of Political Economy* 119 (5): 959–981.

Hill, Ryan and Carolyn Stein. 2020a. “Race to the Bottom: Competition and Quality in Science.” Working Paper. MIT.

— 2020b. “Scooped! Estimating Rewards for Priority in Science.” Working Paper. MIT.

Hira, Ron. 2010. “The H-1B and L-1 Visa Programs: Out of Control.” Briefing Paper 280. Economic Policy Institute.

Holmes, Thomas J. 2011. “The Diffusion of Wal-Mart and Economies of Density.” *Econometrica* 79 (1): 253–302.

Hopenhayn, Hugo A., Gerard Llobet, and Matthew F. Mitchell. 2006. “Rewarding Sequential Innovators: Prizes, Patents, and Buyouts.” *Journal of Political Economy* 114 (6): 1041–1068.

Hopenhayn, Hugo A. and Matthew F. Mitchell. 2001. “Innovation Variety and Patent Breadth.” *RAND Journal of Economics* 32 (1): 152–166.

Hopenhayn, Hugo and Francesco Squintani. 2021. “On the Direction of Innovation.” *Journal of Political Economy* 129 (7): 1991–2022.

Hounshell, David A. and John K. Smith. 1988. *Science and Corporate Strategy: DuPont R&D, 1902-1980*. Cambridge, MA: Cambridge University Press.

Howell, Sabrina T. 2017. “Financing Innovation: Evidence from R&D Grants.” *American Economic Review* 107 (4): 1136–1164.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. 2019. “The Allocation of Talent and U.S. Economic Growth.” *Econometrica* 87 (5): 1439–1474.

Hsieh, Chang-Tai and Peter J. Klenow. 2009. “Misallocation and Manufacturing TFP in China and India.” *Quarterly Journal of Economics* 124 (4): 1403–1448.

Hunt, Jennifer. 2011. “Which Immigrants Are Most Innovative and Entrepreneurial? Distinctions by Entry Visa.” *Journal of Labor Economics* 29 (3): 417–457.

Hunt, Jennifer, Jean-Philippe Garant, Hannah Herman, and David Munroe. 2013. "Why Are Women Underrepresented Amongst Patentees?" *Research Policy* 42 (4): 831–843.

Hunt, Jennifer and Marjolaine Gauthier-Loiselle. 2010. "How Much Does Immigration Boost Innovation?" *American Economic Journal: Macroeconomics* 2 (2): 31–56.

Hvide, Hans K. and Benjamin F. Jones. 2018. "University Innovation and the Professor's Privilege." *American Economic Review* 108 (7): 1860–1898.

Igami, Mitsuru. 2017. "Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981–1998." *Journal of Political Economy* 125 (3): 50.

Igami, Mitsuru and Kosuke Uetake. 2020. "Mergers, Innovation, and Entry-Exit Dynamics: Consolidation of the Hard Disk Drive Industry, 1996–2016." *Review of Economic Studies* 87 (6): 2672–2702.

Jacob, Brian A. and Lars Lefgren. 2011. "The Impact of Research Grant Funding on Scientific Productivity." *Journal of Public Economics* 95 (9): 1168–1177.

Jaffe, Adam B. 1986. "Technological Opportunity and Spillovers of R & D: Evidence from Firms' Patents, Profits, and Market Value." *American Economic Review* 76 (5): 984–1001.

Jaffe, Adam B., Manuel Trajtenberg, and Michael S. Fogarty. 2000. "Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors." *American Economic Review* 90 (2): 215–218.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics* 108 (3): 577–598.

Jaravel, Xavier. 2018. "What Is the Impact of Food Stamps on Prices and Product Variety? The Importance of the Supply Response." *AEA Papers and Proceedings* 108: 557–561.

— 2019. "The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector." *Quarterly Journal of Economics* 134 (2): 715–783.

Jensen, Kyle, Balázs Kovács, and Olav Sorenson. 2018. "Gender Differences in Obtaining and Maintaining Patent Rights." *Nature Biotechnology* 36 (4): 307–309.

Jensen, Kyle and Fiona Murray. 2005. "Intellectual Property Landscape of the Human Genome." *Science* 310 (5746): 239–240.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?" *Review of Economic Studies* 76 (1): 283–317.

— 2010. "Age and Great Invention." *Review of Economics and Statistics* 92 (1): 1–14.

— 2011. "As Science Evolves, How Can Science Policy?" In: *Innovation Policy and the Economy.* Ed. by Josh Lerner and Scott Stern. Vol. 11. NBER Book Series Innovation Policy and the Economy. University of Chicago Press: 103–131.

Jones, Benjamin F. and Lawrence H. Summers. Forthcoming. "A Calculation of the Social Returns to Innovation." In: *Innovation and Public Policy*. Ed. by Benjamin F. Jones and Austan Goolsbee. National Bureau of Economic Research.

Jones, Benjamin F. and Bruce A. Weinberg. 2011. "Age Dynamics in Scientific Creativity." *Proceedings of the National Academy of Sciences of the United States of America* 108 (47): 18910–18914.

Jones, Charles. 2019. "Taxing Top Incomes in a World of Ideas." NBER Working Paper 25725. National Bureau of Economic Research.

Jones, Charles I. 1999. "Growth: With or Without Scale Effects?" *American Economic Review* 89 (2): 139–144.

Jones, Charles I. and Jihee Kim. 2018. "A Schumpeterian Model of Top Income Inequality." *Journal of Political Economy* 126 (5): 1785–1826.

Jones, Charles I. and John Williams. 1998. "Measuring the Social Return to R&D." *Quarterly Journal of Economics* 113 (4): 1119–1135.

Judd, Kenneth L. 2003. "Closed-Loop Equilibrium in a Multi-Stage Innovation Race." *Economic Theory* 21 (2-3): 673–695.

Kamepalli, Sai Krishna, Raghuram G. Rajan, and Luigi Zingales. 2020. "Kill Zone." NBER Working Paper 27146. National Bureau of Economic Research.

Kang, Hyo. 2020. "How Does Competition Affect Innovation? Evidence from U.S. Antitrust Cases." SSRN Scholarly Paper 3516974. Social Science Research Network: 34.

Kantor, Shawn and Alexander Whalley. 2019. "Research Proximity and Productivity: Long-Term Evidence from Agriculture." *Journal of Political Economy* 127 (2): 819–854.

Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73 (1): 39–68.

Keller, Wolfgang. 2002. "Geographic Localization of International Technology Diffusion." *American Economic Review* 92 (1): 120–142.

— 2004. "International Technology Diffusion." *Journal of Economic Literature* 42 (3): 752–782.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. Forthcoming. "Measuring Technological Innovation over the Long Run." *American Economic Review: Insights*.

Kerr, Sari Pekkala, William R. Kerr, and William F. Lincoln. 2015. "Skilled Immigration and the Employment Structures of US Firms." *Journal of Labor Economics* 33 (S1): S147–S186.

Kerr, William R. 2008a. "Ethnic Scientific Communities and International Technology Diffusion." *Review of Economics and Statistics* 90 (3): 518–537.

— 2008b. "The Ethnic Composition of US Inventors." HBS Working Paper 08-006. Harvard Business School.

Kerr, William R. and William F. Lincoln. 2010. “The Supply Side of Innovation: H-1B Visa Reforms and U.S. Ethnic Invention.” *Journal of Labor Economics* 28 (3): 473–508.

Khan, B. Zorina. 2009. *The Democratization of Invention: Patents and Copyrights in American Economic Development, 1790–1920.* 1st edition. Cambridge, United Kingdom: Cambridge University Press.

— 2015. “Inventing Prizes: A Historical Perspective on Innovation Awards and Technology Policy.” *Business History Review* 89 (4): 631–660.

Khan, B. Zorina and Kenneth L. Sokoloff. 1993. “"Schemes of Practical Utility": Entrepreneurship and Innovation Among "Great Inventors" in the United States, 1790-1865.” *Journal of Economic History* 53 (2): 289–307.

Khanna, Gaurav and Munseob Lee. 2020. “High-Skill Immigration, Innovation, and Creative Destruction.” In: *The Roles of Immigrants and Foreign Students in US Science, Innovation, and Entrepreneurship.* Ed. by Ina Ganguli, Shulamit Kahn, and Megan MacGarvie. University of Chicago Press: 73–98.

Kinnunen, Jussi. 1996. “Gabriel Tarde as a Founding Father of Innovation Diffusion Research.” *Acta Sociologica* 39 (4): 431–442.

Kitch, Edmund W. 1977. “The Nature and Function of the Patent System.” *Journal of Law and Economics* 20 (2): 265–290.

Klenow, Peter J. and Andrés Rodríguez-Clare. 1997. “The Neoclassical Revival and Growth Economics: Has It Gone Too Far?” In: *NBER Macroeconomics Annual 1997.* Ed. by Ben S. Bernanke and Julio Rotemberg. Vol. 12. Cambridge, MA: MIT Press: 73–114.

Klette, Tor Jakob and Samuel Kortum. 2004. “Innovating Firms and Aggregate Innovation.” *Journal of Political Economy* 112 (5): 986–1018.

Kline, Patrick, Neviana Petkova, Heidi L. Williams, and Owen Zidar. 2019. “Who Profits from Patents? Rent-Sharing at Innovative Firms.” *Quarterly Journal of Economics* 134 (3): 1343–1404.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. “Technological Innovation, Resource Allocation, and Growth.” *Quarterly Journal of Economics* 132 (2): 665–712.

Kolev, Julian, Yuly Fuentes-Medel, and Fiona Murray. 2020. “Gender Differences in Scientific Communication and Their Impact on Grant Funding Decisions.” *AEA Papers and Proceedings* 110: 245–249.

Koning, Rembrand, Sampsa Samila, and John-Paul Ferguson. 2020. “Inventor Gender and the Direction of Invention.” *AEA Papers and Proceedings* 110: 250–254.

Kremer, Michael. 1998. “Patent Buyouts: A Mechanism for Encouraging Innovation.” *Quarterly Journal of Economics* 113 (4): 1137–1167.

Kremer, Michael and Rachel Glennerster. 2004. *Strong Medicine.* Princeton, NJ: Princeton University Press.

Krugman, Paul. 1991. *Geography and Trade*. Cambridge, MA: MIT Press.

Kügler, Alice. 2021. "The Responsiveness of Inventing: Evidence from a Patent Fee Reform." Job Market Paper.

Kyle, Margaret K. 2006. "The Role of Firm Characteristics in Pharmaceutical Product Launches." *RAND Journal of Economics* 37 (3): 602–618.

— 2007. "Pharmaceutical Price Controls and Entry Strategies." *Review of Economics and Statistics* 89 (1): 88–99.

Kyle, Margaret K. and Anita M. McGahan. 2012. "Investments in Pharmaceuticals Before and After TRIPS." *Review of Economics and Statistics* 94 (4): 1157–1172.

Kyle, Margaret K. and Yi Qian. 2014. "Intellectual Property Rights and Access to Innovation: Evidence from TRIPS." NBER Working Paper 20799. National Bureau of Economic Research.

Lach, Saul and Mark Schankerman. 2008. "Incentives and Invention in Universities." *RAND Journal of Economics* 39 (2): 403–433.

Lakdawalla, Darius and Tomas Philipson. 2012. "Does Intellectual Property Restrict Output? An Analysis of Pharmaceutical Markets." *Journal of Law & Economics* 55 (1): 151–187.

Lamoreaux, Naomi R. and Kenneth L. Sokoloff. 1999. "Inventors, Firms, and the Market for Technology in the Late Nineteenth and Early Twentieth Centuries." In: *Learning by Doing in Markets, Firms, and Countries*. Ed. by Daniel M. G. Raff and Peter Temin. Chicago, IL: University of Chicago Press: 19–60.

Lamoreaux, Naomi R., Kenneth L. Sokoloff, and Dhanoos Sutthiphisal. 2013. "Patent Alchemy: The Market for Technology in US History." *Business History Review* 87 (1): 3–38.

Lampe, Ryan and Petra Moser. 2016. "Patent Pools, Competition, and Innovation—Evidence from 20 US Industries under the New Deal." *Journal of Law, Economics, and Organization* 32 (1): 1–36.

Lanjouw, Jean O. and Mark Schankerman. 2001. "Characteristics of Patent Litigation: A Window on Competition." *RAND Journal of Economics* 32 (1): 129–151.

— 2004. "Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators*." *The Economic Journal* 114 (495): 441–465.

Latour, Bruno and Steve Woolgar. 1986. *Laboratory Life*. Princeton, NJ: Princeton University Press.

Lazonick, William. 1979. "Industrial Relations and Technical Change: The Case of the Self-Acting Mule." *Cambridge Journal of Economics* 3 (3): 231–262.

Le, Trinh and Adam B. Jaffe. 2017. "The Impact of R&D Subsidy on Innovation: Evidence from New Zealand Firms." *Economics of Innovation and New Technology* 26 (5): 429–452.

Lee, Chang-Yang. 2005. "A New Perspective on Industry R&d and Market Structure." *Journal of Industrial Economics* 53 (1): 101–122.

Lee, Tom and Louis L. Wilde. 1980. "Market Structure and Innovation: A Reformulation." *Quarterly Journal of Economics* 94 (2): 429–436.

Leibenstein, Harvey. 1966. "Allocative Efficiency vs. "X-Efficiency"." *American Economic Review* 56 (3): 392–415.

Lemley, Mark A. 2001. "Rational Ignorance at the Patent Office." *Northwestern University Law Review* 95 (4): 1–34.

Lemley, Mark A. and Bhaven N. Sampat. 2008. "Is the Patent Office a Rubber Stamp?" *Emory Law Journal* 58: 101–128.

Lemley, Mark A. and Carl Shapiro. 2005. "Probabilistic Patents." *Journal of Economic Perspectives* 19 (2): 75–98.

Lerner, Josh. 1994. "The Importance of Patent Scope: An Empirical Analysis." *RAND Journal of Economics* 25 (2): 319–333.

— 1999. "The Government as Venture Capitalist: The Long-Run Impact of the SBIR Program." *Journal of Business* 72 (3): 285–318.

— 2009. "The Empirical Impact of Intellectual Property Rights on Innovation: Puzzles and Clues." *American Economic Review* 99 (2): 343–348.

— 2012. *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed – and What to Do about It*. The Kauffman Foundation Series on Innovation and Entrepreneurship. Princeton, NJ: Princeton University Press.

Lerner, Josh and Amit Seru. 2017. "The Use and Misuse of Patent Data: Issues for Corporate Finance and Beyond." NBER Working Paper 24053. National Bureau of Economic Research.

Lerner, Josh and Jean Tirole. 2006. "A Model of Forum Shopping." *American Economic Review* 96 (4): 1091–1113.

— 2015. "Standard-Essential Patents." *Journal of Political Economy* 123 (3): 547–586.

Letina, Igor, Armin Schmutzler, and Regina Seibel. 2021. "Killer Acquisitions and Beyond: Policy Effects on Innovation Strategies." CEPR Discussion Paper 15167. Centre for Economic Policy Research.

Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, Sidney G. Winter, Richard J. Gilbert, and Zvi Griliches. 1987. "Appropriating the Returns from Industrial Research and Development." *Brookings Papers on Economic Activity* 3 (1): 783–831.

Li, Danielle. 2017. "Expertise versus Bias in Evaluation: Evidence from the NIH." *American Economic Journal: Applied Economics* 9 (2): 60–92.

Li, Danielle and Leila Agha. 2015. "Big Names or Big Ideas: Do Peer-Review Panels Select the Best Science Proposals?" *Science* 348 (6233): 434–438.

Li, Xing, Megan MacGarvie, and Petra Moser. 2018. "Dead Poets' Property—How Does Copyright Influence Price?" *RAND Journal of Economics* 49 (1): 181–205.

Liebowitz, S. J. and Stephen E. Margolis. 1995. "Path Dependence, Lock-in, and History." *Journal of Law, Economics, & Organization* 11 (1): 205–226.

Liu, Lu, Yang Wang, Roberta Sinatra, C. Lee Giles, Chaoming Song, and Dashun Wang. 2018. "Hot Streaks in Artistic, Cultural, and Scientific Careers." *Nature* 559 (7714): 396–399.

Lobel, Orly. 2019. *You Don't Own Me: The Court Battles That Exposed Barbie's Dark Side*. New York, NY: WW Norton.

Long, Genia. 2019. "Federal Government-Interest Patent Disclosures for Recent Top-Selling Drugs." *Journal of Medical Economics* 22 (12): 1261–1267.

Loury, Glenn C. 1979. "Market Structure and Innovation." *Quarterly Journal of Economics* 93 (3): 395–410.

Lück, Sonja, Benjamin Balsmeier, Florian Seliger, and Lee Fleming. 2020. "Early Disclosure of Invention and Reduced Duplication: An Empirical Test." *Management Science* 66 (6): 2677–2685.

Lucking, Brian, Nicholas Bloom, and John Van Reenen. 2018. "Have R&D Spillovers Changed?" NBER Working Paper 24622. National Bureau of Economic Research.

Machlup, Fritz. 1962. *The Production and Distribution of Knowledge in the United States*. Princeton, NJ: Princeton University Press.

Machlup, Fritz and Edith Penrose. 1950. "The Patent Controversy in the Nineteenth Century." *Journal of Economic History* 10 (1): 1–29.

Maclaurin, W. Rupert. 1950. "The Process of Technological Innovation: The Launching of a New Scientific Industry." *American Economic Review* 40 (1): 90–112.

— 1953. "The Sequence from Invention to Innovation and Its Relation to Economic Growth." *Quarterly Journal of Economics* 67 (1): 97–111.

Mankiw, N. Gregory and Michael D. Whinston. 1986. "Free Entry and Social Inefficiency." *RAND Journal of Economics* 17 (1): 48.

Mansfield, Edwin. 1982. "Tax Policy and Innovation." *Science* 215 (4538): 1365–1371.

— 1986a. "Patents and Innovation: An Empirical Study." *Management Science* 32 (2): 173–181.

— 1986b. "The R&D Tax Credit and Other Technology Policy Issues." *American Economic Review* 76 (2): 190–194.

Manso, Gustavo. 2011. "Motivating Innovation." *Journal of Finance* 66 (5): 1823–1860.

Manton, Kenneth G., Xi-Liang Gu, Gene Lowrimore, Arthur Ullian, and H. Dennis Tolley. 2009. "NIH Funding Trajectories and Their Correlations with US Health Dynamics from 1950 to 2004." *Proceedings of the National Academy of Sciences* 106 (27): 10981–10986.

Marco, Alan C., Amanda Myers, Stuart J. H. Graham, Paul D'Agostino, and Kirsten Apple. 2015. "The USPTO Patent Assignment Dataset: Descriptions and Analysis." USPTO Economic Working Paper 2015-2.

Marschke, Gerald, Allison Nunez, Bruce A. Weinberg, and Huifeng Yu. 2018. "Last Place? The Intersection of Ethnicity, Gender, and Race in Biomedical Authorship." *AEA Papers and Proceedings* 108: 222–227.

Marshall, Alfred. 1890. *Principles of Economics*. Vol. I. London: MacMillan & Co.

Marshall, Guillermo and Álvaro Parra. 2019. "Innovation and Competition: The Role of the Product Market." *International Journal of Industrial Organization* 65: 221–247.

Marx, Matt and Lee Fleming. 2012. "Non-Compete Agreements: Barriers to Entry ... and Exit?" *Innovation Policy and the Economy* 12 (1): 39–64.

Marx, Matt and Aaron Fuegi. 2020. "Reliance on Science by Inventors: Hybrid Extraction of In-Text Patent-to-Article Citations." NBER Working Paper 27987. National Bureau of Economic Research.

Marx, Matt, Deborah Strumsky, and Lee Fleming. 2009. "Mobility, Skills, and the Michigan Non-Compete Experiment." *Management Science* 55 (6): 875–889.

Matloff, Norman. 2003. "On the Need for Reform of the H-1B Non-Immigrant Work Visa in Computer-Related Occupations." *University of Michigan Journal of Law Reform* 36 (4): 815–913.

McElheran, Kristina, Nathan Goldschlag, Zachary Kroff, David Beede, Lucia Foster, Erik Brynjolfsson, Catherine Buffington, and Emin Dinlersoz. 2020. "Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey." Conference Paper.

Merton, Robert K. 1935. "Fluctuations in the Rate of Industrial Invention." *Quarterly Journal of Economics* 49 (3): 454–474.

— 1957. "Priorities in Scientific Discovery: A Chapter in the Sociology of Science." *American Sociological Review* 22 (6): 635–659.

— 1968. "The Matthew Effect in Science: The Reward and Communication Systems of Science Are Considered." *Science* 159 (3810): 56–63.

— 1969. "Behavior Patterns of Science." *American Scientist* 57 (1): 1–23.

— 1973. *The Sociology of Science: Theoretical and Empirical Incvestigations*. Chicago, IL: University of Chicago Press.

— 1988. "The Matthew Effect in Science, II: Cumulative Advantage and the Symbolism of Intellectual Property." *Isis* 79 (4): 606–623.

Miller, Chad R. and Brian Richard. 2010. "The Policy Diffusion of the State R&D Investment Tax Credit." *Journal of Economic Perspectives* 42 (1): 22–35.

Miller, Shawn P. 2013. "What's the Connection between Repeat Litigation and Patent Quality? A (Partial) Defense of the Most Litigated Patents." *Stanford Technology Law Review* 16 (2): 313–348.

Mokyr, Joel. 1999. "Editor's Introduction: The New Economic History and the Industrial Revoluition." In: *The British Industrial Revolution*. New York, NY: Routledge: 127.

— 2010. "The Contribution of Economic History to the Study of Innovation and Technical Change 1750–1914." In: *Handbook of The Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. North-Holland: 11–50.

Mokyr, Joel, Chris Vickers, and Nicolas L. Ziebarth. 2015. "The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different?" *Journal of Economic Perspectives* 29 (3): 31–50.

Moore, Kimberly A. 2005. "Worthless Patents." *Berkeley Technology Law Journal* 20 (4): 1521–1552.

Moretti, Enrico, Claudia Steinwender, and John Van Reenen. 2019. "The Intellectual Spoils of War? Defense R&D, Productivity and International Spillovers." NBER Working Paper 26483. National Bureau of Economic Research.

Moretti, Enrico and Daniel J. Wilson. 2014. "State Incentives for Innovation, Star Scientists and Jobs: Evidence from Biotech." *Journal of Urban Economics* 79: 20–38.

— 2017. "The Effect of State Taxes on the Geographical Location of Top Earners: Evidence from Star Scientists." *American Economic Review* 107 (7): 1858–1903.

Moscona, Jacob. 2020. "Flowers of Invention: Patent Protection and Productivity Growth in US Agriculture." Working Paper. MIT.

Moser, Petra. 2005. "How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs." *American Economic Review* 95 (4): 1214–1236.

— 2012. "Innovation without Patents: Evidence from World's Fairs." *Journal of Law and Economics* 55 (1): 43–74.

Moser, Petra, Kamran Bilir, and Irina Talis. 2011. "Do Treaties Encourage Technology Transfer? Evidence from the Paris Convention." SSRN Scholarly Paper ID 1893052. Rochester, NY: Social Science Research Network.

Moser, Petra, Joerg Ohmstedt, and Paul W. Rhode. 2018. "Patent Citations and the Size of the Inventive Step - Evidence from Hybrid Corn." *Management Science* 64 (4): 1477–1973.

Moser, Petra and Samuel San. 2019. "Immigration, Science, and Invention: Evidence from the Quota Acts." SSRN Scholarly Paper 3558718. Social Science Research Network.

Moser, Petra, Alessandra Voena, and Fabian Waldinger. 2014. "German Jewish Émigrés and US Invention." *American Economic Review* 104 (10): 3222–3255.

Mowery, David C., Richard R. Nelson, Bhaven N. Sampat, and Arvids A. Ziedonis. 2015. *Ivory Tower and Industrial Innovation: University-Industry Technology Transfer Before and After the Bayh-Dole Act*. Palo Alto, CA: Stanford University Press.

Mowery, David C. and Nathan Rosenberg. 1991. "The Commercial Exploitation of Science." In: *Technology and the Pursuit of Economic Growth*. Cambridge, United Kingdom: Cambridge University Press: 45–54.

Murray, Fiona, Philippe Aghion, Mathias Dewatripont, Julian Kolev, and Scott Stern. 2016. "Of Mice and Academics: Examining the Effect of Openness on Innovation." *American Economic Journal: Economic Policy* 8 (1): 212–252.

Murray, Fiona and Scott Stern. 2007. "Do Formal Intellectual Property Rights Hinder the Free Flow of Scientific Knowledge? An Empirical Test of the Anti-Commons Hypothesis." *Journal of Economic Behavior and Organization* 63 (4): 648–687.

Myers, Kyle. 2020. "The Elasticity of Science." *American Economic Journal: Applied Economics* 12 (4): 103–134.

Myers, Kyle and Lauren Lanahan. 2021. "Estimating Spillovers from Publicly Funded R&D: Evidence from the US Department of Energy." SSRN Scholarly Paper ID 3550479. Rochester, NY: Social Science Research Network.

Nagaraj, Abhishek and Scott Stern. 2020. "The Economics of Maps." *Journal of Economic Perspectives* 34 (1): 196–221.

National Bureau of Economic Research. 1962. *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press.

National Science Foundation. 2018. "Science & Engineering Indicators 2018." Report NSB-2018-1. Alexandria, VA: National Science Foundation.

Nelson, Richard R. 1959a. "The Economics of Invention: A Survey of the Literature." *Journal of Business* 32 (2): 101–127.

— 1959b. "The Simple Economics of Basic Scientific Research." *Journal of Political Economy* 67 (3): 297–306.

— 1962. "Introduction to "The Rate and Direction of Inventive Activity: Economic and Social Factors"." In: *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press: 1–16.

Nelson, Richard R. and Edmund S. Phelps. 1966. "Investment in Humans, Technological Diffusion, and Economic Growth." *American Economic Review* 56 (1-2): 69–75.

Nelson, Richard R. and Sidney G. Winter. 1985. *An Evolutionary Theory of Economic Change*. Cambridge, MA.: Belknap Press: An Imprint of Harvard University Press.

Newham, Melissa and Rune Midjord. 2018. "Herd Behavior in FDA Committees: A Structural Approach." DIW Discussion Paper 1744. DIW Berlin.

Nicholas, Tom. 2011. "Cheaper Patents." *Research Policy* 40 (2): 325–339.

Nisbet, Richard. 1980. *History of the Idea of Progress*. New York, NY: Basic Books.

Nordhaus, William D. 1969. *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*. Cambridge, MA: MIT Press.

Nuvolari, Alessandro and James Sumner. 2013. "Inventors, Patents, and Inventive Activities in the English Brewing Industry, 1634–1850." *Business History Review* 87 (1): 95–120.

Ogburn, William Fielding. 1922. *Social Change with Respect to Culture and Original Nature*. New York, NY: B.W. Huebsch, Inc.

Olmstead-Rumsey, Jane. 2021. "Market Concentration and the Productivity Slowdown." Working Paper.

Organization for Economic Co-operation and Development. 2020. "Measuring Tax Support for R&D and Innovation." Technical Report. OECD.

Ottaviani, Marco. 2020. "Grantmaking." CEPR Discussion Paper 15389. Centre for Economic Policy Research.

Ouellette, Lisa Larrimore. 2012. "Do Patents Disclose Useful Information?" *Harvard Journal of Law & Technology* 25 (2): 531–593.

Ouellette, Lisa Larrimore and Andrew Tutt. 2020. "How Do Patent Incentives Affect University Researchers?" *International Review of Law and Economics* 61: Article 105883.

Pakes, Ariel. 1986. "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks." *Econometrica* 54 (4): 755–784.

Peirce, C. S. 1879. "Note on the Theory of the Economy of Research." *United States Coast Survey*. Reprinted in 1967 in *Operations Research* 15 (4): 642-8.

Pekari, Ari-Pekka. 1993. "Työsuhdekeksintöjen Merkitys Ja Työsuhdekeksintöjen Syntumiseen Vaikuttavat Tekijät Suomalaisissa Teollisuusyrityksissä." (The significance of employee inventions and factors affecting employee inventions in Finnish industrial firms; in Finnish only). Sitra.

Peri, Giovanni, Kevin Shih, and Chad Sparber. 2015. "STEM Workers, H-1B Visas, and Productivity in US Cities." *Journal of Labor Economics* 33 (S1): S225–S255.

Pezzoni, Michele, Reinhilde Veugelers, and Fabiana Visentin. 2019. "How Fast Is a Novel Technology Going to Be a Hit? Antecedents Predicting Follow-on Inventions." *Academy of Management Proceedings* 2019 (1): Article 13223.

Phillips, Gordon M. and Alexei Zhdanov. 2013. "R&D and the Incentives from Merger and Acquisition Activity." *Review of Financial Studies* 26 (1): 34–78.

Piketty, Thomas and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913-1998." *Quarterly Journal of Economics* 118 (1): 1–39.

Plant, Arnold. 1934. "The Economic Theory Concerning Patents for Inventions." *Economica* 1 (1): 30–51.

Pless, Jacquelyn. 2021. "Are 'Complementary Policies' Substitutes? Evidence from R&D Subsidies in the UK." Tech. rep. Job Market Paper.

Polanvyi, Michael. 1944. "Patent Reform." *Review of Economic Studies* 11 (2): 61–76.

Rai, Arti K. and Bhaven N. Sampat. 2012. "Accountability in Patenting of Federally Funded Research." *Nature Biotechnology* 30 (10): 953–956.

Rasmusen, Eric. 1988. "Entry for Buyout." *Journal of Industrial Economics* 36 (3): 281–299.

Reinganum, Jennifer F. 1982. "A Dynamic Game of R and D: Patent Protection and Competitive Behavior." *Econometrica* 50 (3): 671–688.

— 1985. "Innovation and Industry Evolution." *Quarterly Journal of Economics* 100 (1): 81–99.

Richter, Ralf and Jochen Streb. 2011. "Catching-Up and Falling Behind: Knowledge Spillover from American to German Machine Toolmakers." *The Journal of Economic History* 71 (4): 1006–1031.

Righi, Cesare and Timothy Simcoe. 2019. "Patent Examiner Specialization." *Research Policy* 48 (1): 137–148.

Roach, Michael and Wesley M. Cohen. 2013. "Lens or Prism? Patent Citations as a Measure of Knowledge Flows from Public Research." *Management Science* 59 (2): 504–525.

Rockett, Katharine, Bronwyn H. Hall, and Nathan Rosenberg. 2010. "Property Rights and Invention." In: *Handbook of the Economics of Innovation*. Vol. 1. Elsevier: 315–380.

Rogers, Everett M. 1962. *Diffusion of Innovations*. New York, NY: Free Press of Glencoe.

Roin, Benjamin N. 2005. "The Disclosure Function of the Patent System (Or Lack Thereof)." *Harvard Law Review* 118 (6): 2007–2028.

— 2014a. "Intellectual Property versus Prizes: Reframing the Debate." *University of Chicago Law Review* 81 (3).

— 2014b. "The Case for Tailoring Patent Awards Based on Time-to-Market." *UCLA Law Review* 61: 672–759.

Romer, Paul M. 1986. "Increasing Returns and Long-Run Growth." *Journal of Political Economy* 94 (5): 1002–1037.

— 1990. "Endogenous Technological Change." *Journal of Political Economy* 98 (5): S71–S102.

Rosenberg, Nathan. 1990. "Why Do Firms Do Basic Research (with Their Own Money)?" *Research Policy* 19 (2): 165–174.

Roy, A. D. 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers* 3 (2): 135–146.

— 1952. "Quota Restriction and Goldbricking in a Machine Shop." *American Journal of Sociology* 57 (5): 427–442.

Ruttan, Vernon W. 1959. "Usher and Schumpeter on Invention, Innovation, and Technological Change." *Quarterly Journal of Economics* 73 (4): 596–606.

Ryan, Bryce and Neal Gross. 1943. "The Diffusion of Hybrid Seed Corn In Two Iowa Communities." *Rural Sociology* 8 (1): 15–24.

Sakakibara, Mariko and Lee G. Branstetter. 2001. "Do Stronger Patents Induce More Innovation? Evidence from the 1988 Japanese Patent Law Reforms." *RAND Journal of Economics* 32 (1): 77–100.

Sampat, Bhaven N. 2015. “Serendipity.” SSRN Scholarly Paper ID 2545515. Rochester, NY: Social Science Research Network.

Sampat, Bhaven N. and Frank R. Lichtenberg. 2011. “What Are The Respective Roles Of The Public And Private Sectors In Pharmaceutical Innovation?” *Health Affairs* 30 (2): 332–339.

Sampat, Bhaven N. and Heidi L. Williams. 2019. “How Do Patents Affect Follow-On Innovation? Evidence from the Human Genome.” *American Economic Review* 109 (1): 203–236.

Samuelson, Paul A. 1989. “Ricardo Was Right!” *Scandinavian Journal of Economics* 91 (1): 47–62.

Santoleri, Pietro, Andrea Mina, Alberto Di Minin, and Irene Martelli. 2020. “The Causal Effects of R&D Grants: Evidence from a Regression Discontinuity.” LEM Working Paper 2020/18. Laboratory of Economics and Management.

Sarada, Sarada, Michael J. Andrews, and Nicolas L. Ziebarth. 2019. “Changes in the Demographics of American Inventors, 1870–1940.” *Explorations in Economic History* 74: 101275.

Schaffer, Simon. 1994. “Making Up Discovery.” In: *Dimensions of Creativity*. Ed. by Margaret A. Boden. Cambridge, MA: MIT Press: 13–51.

Schankerman, Mark and Ariel Pakes. 1986. “Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period.” *Economic Journal* 96 (384): 1052–1076.

Schankerman, Mark and Florian Schuett. Forthcoming. “Screening for Patent Quality: Examination, Fees, and the Courts.” *Review of Economic Studies*.

Scherer, Frederic M. 1965. “Corporate Inventive Output, Profits, and Growth.” *Journal of Political Economy* 73 (3): 290–297.

— 1967. “Market Structure and the Employment of Scientists and Engineers.” *American Economic Review* 57 (3): 524–531.

— 1972. “Nordhaus’ Theory of Optimal Patent Life: A Geometric Reinterpretation.” *American Economic Review* 62 (3): 422–427.

— 1982. “Demand-Pull and Technological Invention: Schmookler Revisted.” *Journal of Industrial Economics* 30 (3): 225–237.

Schmitz Jr., James A. 2005. “What Determines Productivity? Lessons from the Dramatic Recovery of the U.S. and Canadian Iron Ore Industries Following Their Early 1980s Crisis.” *Journal of Political Economy* 113 (3): 582–625.

Schmookler, Jacob. 1954. “The Level of Inventive Activity.” *Review of Economics and Statistics* 36 (2): 183–190.

— 1962. “Economic Sources of Inventive Activity.” *Journal of Economic History* 22 (1): 1–20.

— 1966. *Invention and Economic Growth*. Cambridge, MA: Harvard University Press.

Schumpeter, Joseph A. 1934. *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*. New York, NY: Oxford University Press.

Schumpeter, Joseph A. 1939. *Business Cycles: A Theoretical and Statistical Analysis of the Capitalist Process.* New York, NY: McGraw-Hill Book Company.

— 1942. *Capitalism, Socialism and Democracy.* New York, NY: Harper & Brothers.

Scotchmer, Suzanne. 1999. "On the Optimality of the Patent Renewal System." *RAND Journal of Economics* 30 (2): 181–196.

— 2004. "The Political Economy of Intellectual Property Treaties." *Journal of Law, Economics, and Organization* 20 (2): 415–437.

Scotchmer, Suzanne and Jerry Green. 1990. "Novelty and Disclosure in Patent Law." *RAND Journal of Economics* 21 (1): 131–146.

Segal, Ilya and Michael D. Whinston. 2007. "Antitrust in Innovative Industries." *American Economic Review* 97 (5): 1703–1730.

Serrano, Carlos J. 2010. "The Dynamics of the Transfer and Renewal of Patents." *RAND Journal of Economics* 41 (4): 686–708.

Shapiro, Carl. 2012. "Competition and Innovation: Did Arrow Hit the Bull's Eye?" In: *The Rate and Direction of Inventive Activity Revisited.* Chicago, IL: University of Chicago Press: 361–404.

Shavell, Steven and Tanguy van Ypersele. 2001. "Rewards versus Intellectual Property Rights." *Journal of Law and Economics* 44 (2): 525–547.

Sheff, Jeremy. 2020. "Jefferson's Taper." *Southern Methodist University Law Review* 73 (2): 299.

Shwed, Uri and Peter S. Bearman. 2010. "The Temporal Structure of Scientific Consensus Formation." *American Sociological Review* 75 (6): 817–840.

Simcoe, Timothy. 2012. "Standard Setting Committees: Consensus Governance for Shared Technology Platforms." *American Economic Review* 102 (1): 305–336.

Skinner, Jonathan and Douglas Staiger. 2015. "Technology Diffusion and Productivity Growth in Health Care." *Review of Economics and Statistics* 97 (5): 951–964.

Slattery, Cailin. 2020. "Bidding for Firms." Columbia Business School.

Soete, Luc, Bart Verspagen, and Bas ter Weel. 2010. "Systems of Innovation." In: *Handbook of the Economics of Innovation.* Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 2. North-Holland: 1159–1180.

Solo, Carolyn Shaw. 1951. "Innovation in the Capitalist Process: A Critique of the Schumpeterian Theory." *Quarterly Journal of Economics* 65 (3): 417–428.

Solow, Robert M. 1956. "A Contribution to the Theory of Economic Growth." *Quarterly Journal of Economics* 70 (1): 65.

— 1957. "Technical Change and the Aggregate Production Function." *Review of Economics and Statistics* 39 (3): 312–320.

Stafford, Alfred B. 1952. “Is the Rate of Invention Declining?” *American Journal of Sociology* 57 (6): 539–545.

Stamp, Josiah. 1929. *Some Economic Factors in Modern Life*. London: P.S. King & Son, Ltd.

Stephan, Paula E. 2010. “The Economics of Science.” In: *Handbook of the Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 1. Elsevier: 217–273.

— 2015a. *How Economics Shapes Science*. Cambridge, MA: Harvard University Press.

— 2015b. “The Endless Frontier: Reaping What Bush Sowed?” In: *The Changing Frontier: Rethinking Science and Innovation Policy*. Chicago, IL: University of Chicago Press: 321–366.

Stephan, Paula E. and Sharon G. Levin. 2001. “Exceptional Contributions to US Science by the Foreign-Born and Foreign-Educated.” *Population Research and Policy Review* 20 (1): 59–79.

Stern, Scott. 2004. “Do Scientists Pay to Be Scientists?” *Management Science* 50 (6): 835–853.

Stokes, Donald E. 1997. *Pasteur’s Quadrant: Basic Science and Technological Innovation*. Brookings Institution Press.

Stoneman, Paul and Giuliana Battisti. 2010. “The Diffusion of New Technology.” In: *Handbook of the Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 2. Elsevier: 733–760.

Sutch, Richard, Gary D. Libecap, and Richard H. Steckel. 2011. “The Impact of the 1936 Corn Belt Drought on American Farmers’ Adoption of Hybrid Corn.” In: *The Economics of Climate Change: Adaptations Past and Present*. Chicago, IL: University of Chicago Press: 195–223.

Sutton, John. 1998. *Technology and Market Structure*. Cambridge, MA: MIT Press.

Swan, Trevor W. 1956. “Economic Growth and Capital Accumulation.” *Economic Record* 32 (2): 334–361.

Syverson, Chad. 2004. “Product Substitutability and Productivity Dispersion.” *Review of Economics and Statistics* 86 (2): 534–550.

— 2011. “What Determines Productivity?” *Journal of Economic Literature* 49 (2): 326–365.

Tarde, Gabriel. 1890. *Les Lois de l’imitation*. Paris, France: Félix Alcan.

Teece, David J. 1977. “Technology Transfer by Multinational Firms: The Resource Cost of Transferring Technological Know-How.” *Economic Journal* 87 (346): 242–261.

— 1986. “Profiting from Technological Innovation: Implications for Integration, Collaboration, Licensing and Public Policy.” *Research Policy* 15 (6): 285–305.

— 2010. “Technological Innovation and the Theory of the Firm: The Role of Enterprise-Level Knowledge, Complementarities, and (Dynamic) Capabilities.” In: *Handbook of the Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 1. North-Holland: 679–730.

Terborgh, George W. 1950. “Capitalism and Innovation.” *American Economic Review* 40 (2): 118–123.

Thompson, Peter and Melanie Fox-Kean. 2005. "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment." *American Economic Review* 95 (1): 450–460.

Tirole, Jean. 2020. "Competition and the Industrial Challenge for the Digital Age." Working Paper.

Toivanen, Otto and Lotta Väänänen. 2012. "Returns to Inventors." *Review of Economics and Statistics* 94 (4): 1173–1190.

Toole, Andrew A. 2012. "The Impact of Public Basic Research on Industrial Innovation: Evidence from the Pharmaceutical Industry." *Research Policy* 41 (1): 1–12.

Toole, Andrew A., Michelle J. Saksena, Charles A. W. deGrazia, Katherine P. Black, Francesco Lissoni, Ernest Miguelez, and Gianluca Tarasconi. 2020. "Progress and Potential: 2020 Update on U.S. Women Inventor-Patentees." IP Data Highlights 4. USPTO: 12.

Trajtenberg, Manuel. 1990. "A Penny for Your Quotes: Patent Citations and the Value of Innovations." *RAND Journal of Economics* 21 (1): 172–187.

US Department of State. 2019. "Nonimmigrant Visa Statistics." Annual Report. Bureau of Consular Affairs.

US Department of the Treasury. 2016. "Research and Experimentation (R&E) Credit." Office of Tax Analysis Report. US Department of the Treasury: 22.

US Patent and Trademark Office. 2015. "Patent Litigation and USPTO Trials: Implications for Patent Examination Quality." Technical Report. USPTO.

Usher, Abbott Payson. 1954. *A History of Mechanical Invention: Revised Edition*. Cambridge, MA: Harvard University Press.

USPTO Patent Technology Monitoring Team. 2021. *U. S. Patent Statistics Chart Calendar Years 1963-2020*. https://www.uspto.gov/web/offices/ac/ido/oeip/taf/us_stat.htm.

Uzzi, Brian, Satyam Mukherjee, Michael Stringer, and Ben Jones. 2013. "Atypical Combinations and Scientific Impact." *Science* 342 (6157): 468–472.

Valiki, Keyvan, Florenta Teodoridis, and Michael Bikard. 2020. "Detrimental Collaborations in Creative Work: Evidence from Economics." Working Paper.

Van Reenen, John. 1996. "The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies." *Quarterly Journal of Economics* 111 (1): 195–226.

von Hippel, Eric. 1988. *Sources of Innovation*. New York, NY: Oxford University Press.

— 2010. "Open User Innovation." In: *Handbook of the Economics of Innovation*. Ed. by Bronwyn H. Hall and Nathan Rosenberg. Vol. 1. North-Holland: 411–427.

Waldinger, Fabian. 2016. "Bombs, Brains, and Science: The Role of Human and Physical Capital for the Creation of Scientific Knowledge." *Review of Economics and Statistics* 98 (5): 811–831.

Walsh, John P., Charlene Cho, and Wesley M. Cohen. 2005. "View from the Bench: Patents and Material Transfers." *Science* 309 (5743): 2002–2003.

Walsh, John P., Wesley M. Cohen, and Charlene Cho. 2007. "Where Excludability Matters: Material versus Intellectual Property in Academic Biomedical Research." *Research Policy* 36 (8): 1184–1203.

Watzinger, Martin, Thomas A. Fackler, Markus Nagler, and Monika Schnitzer. 2020. "How Antitrust Enforcement Can Spur Innovation: Bell Labs and the 1956 Consent Decree." *American Economic Journal: Economic Policy* 12 (4): 328–359.

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

Weyl, E. Glen and Jean Tirole. 2012. "Market Power Screens Willingness-to-Pay." *Quarterly Journal of Economics* 127 (4): 1971–2003.

Whittington, Kjersten B. 2009. "Patterns of Male and Female Scientific Dissemination in Public and Private Science." In: *Science and Engineering Careers in the United States: An Analysis of Markets and Employment*. Ed. by Richard B. Freeman and Daniel L. Goroff. Chicago, IL: University of Chicago Press: 194–228.

Williams, Heidi L. 2013. "Intellectual Property Rights and Innovation: Evidence from the Human Genome." *Journal of Political Economy* 121 (1): 1–27.

— 2017. "How Do Patents Affect Research Investments?" *Annual Review of Economics* 9 (1): 441–469.

Wilson, Daniel J. 2009. "Beggar Thy Neighbor? The In-State, Out-of-State, and Aggregate Effects of R&D Tax Credits." *Review of Economics and Statistics* 91 (2): 431–436.

Wright, Brian D. 1983. "The Economics of Invention Incentives: Patents, Prizes, and Research Contracts." *American Economic Review* 73 (4): 691–707.

Wu, Lingfei, Dashun Wang, and James A. Evans. 2019. "Large Teams Develop and Small Teams Disrupt Science and Technology." *Nature* 566 (7744): 378–382.

Wuchty, Stefan, Benjamin F. Jones, and Brian Uzzi. 2007. "The Increasing Dominance of Teams in Production of Knowledge." *Science* 316 (5827): 1036–1039.

Yang, Chenyu. 2020. "Vertical Structure and Innovation: A Study of the SoC and Smartphone Industries." *RAND Journal of Economics* 51 (3): 739–785.

Yin, Wesley. 2008. "Market Incentives and Pharmaceutical Innovation." *Journal of Health Economics* 27 (4): 1060–1077.

— 2009. "R&D Policy, Agency Costs and Innovation in Personalized Medicine." *Journal of Health Economics* 28 (5): 950–962.

Young, H. Peyton. 2009. "Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning." *American Economic Review* 99 (5): 1899–1924.

Zhang, Xiaoquan (Michael) and Feng Zhu. 2011. "Group Size and Incentives to Contribute: A Natural Experiment at Chinese Wikipedia." *American Economic Review* 101 (4): 1601–1615.

Zingales, Luigi, Guy Rolnik, and Filippo Maria Lancieri. 2019. "Stigler Committee on Digital Platforms." Final Report. University of Chicago Booth School of Business: Stigler Center for the Study of the Economy and the State: 336.