

Taxation and Innovation: What Do We Know?*

Ufuk Akcigit and Stefanie Stantcheva

April 29, 2020

Abstract

Tax policies are a wide array of tools, commonly used by governments to influence the economy. In this paper, we review the many margins through which tax policies can affect innovation, the main driver of economic growth in the long-run. These margins include the impact of tax policy on i) the quantity and quality of innovation; ii) the geographic mobility of innovation and inventors across U.S. states and countries; iii) the declining business dynamism in the U.S., firm entry, and productivity; iv) the quality composition of firms, inventors, and teams; and v) the direction of research effort, e.g., toward applied versus basic research, or toward dirty versus clean technologies. We give ideas drawn from research on how the design of policy can allow policy makers to foster the most productive firms without wasting public funds on less productive ones.

1 Introduction

There is a myriad of reasons why we should care about innovation. Innovation is the source of technological progress and the main driver of economic growth in the long-run. In recent work, ? show that U.S. states with the most innovations also witnessed the fastest growth between 1900 and 2000. Beyond its important role in growth, innovation is also strongly associated with social mobility, especially when it is done by new entrants (??), and even with the well-being of people (?).

It is therefore evident why policy makers would try to understand how policies impact innovation and what policy tools can be used to foster it. This issue is particularly pressing in the U.S., as business dynamism has been slowing in the last several decades. Recent studies have documented the many faces of this decline: a lower entry rate of new businesses, a slow-down in productivity growth, a falling labor share in output, and rises in market concentration and the corporate profit share. In this context, tax policy can be a powerful tool. Used correctly, it can provide effective incentives for many economic activities and innovation is no exception. Used inefficiently, it can create heavy deadweight burdens, hurt incentives, and slow down innovation. It is thus critical to innovation to implement the appropriate tax policy.

In this chapter, we will discuss the various roles of tax policy for innovation and ways in which it could be used to foster technological progress at low fiscal cost. When it comes to innovation,

*Akcigit: University of Chicago, CEPR, and NBER (e-mail: uakcigit@uchicago.edu). Stantcheva: Harvard, CEPR, and NBER (e-mail: sstantcheva@fas.harvard.edu). We thank Austan Goolsbee and Ben Jones for very valuable feedback. The authors thank National Science Foundation for financial support.

tax policies can be classified into two broad groups: general tax policy – such as the personal or corporate income tax – and targeted tax policies – such as R&D tax credits, local tax incentives for innovating firms, or subsidies for specific types of research.

In the public imagination, innovation is often viewed as a mysterious process whereby wonderful new things are created almost magically. When we think of path-breaking superstar inventors from history, such as Thomas Edison, Alexander Bell, or Nikola Tesla, the picture that comes to mind is one of hard-working and enthusiastic scientists, who neglect financial incentives and only strive for intellectual achievement. But innovation is an economic activity and the result of intentional effort and investments. It may certainly have a different time profile and shape of risk and returns than other activities. People may also have varying degrees of other motivations – such as social prestige or the love of science– as is the case for other types of economic activities. How strongly innovation responds to economic incentives is ultimately an empirical question.

General taxes are typically set for the purpose of raising revenues and redistributing income; they are typically not set with innovation in mind. Yet they reduce the expected net returns to innovation inputs and can lead to less innovation as an unwelcome by-product. This is an efficiency cost that needs to be taken into account, together with other more standard margins that are considered when setting tax policy (such as labor supply or tax avoidance). The estimates of these efficiency costs in terms of lost innovation could lead to a reassessment of what the right level of taxes should, and would, be an input into our optimal tax formulas (?). More specific tax policies targeted to innovations go a step further and can be designed intentionally so as to foster innovation. It is important to understand all the margins along which they can play a role, as innovation is a complex process made of many steps.

In this chapter, we will provide a conceptual framework to think about the effects of general and targeted tax policy on innovation. A key consideration is that there are many channels and margins through which innovation will respond to tax policies. We outline them in Section 2 and summarize them visually in Figure 1. We then dig into the recent literature that sheds light on each of these channels and response margins. The organization of the paper is as follows. Each section presents the key issue on the margin under consideration and draws out the implications for tax policy.¹ The literature reviewed is by no means comprehensive. Instead, we focus on work that we have done with co-authors and focus on extracting the implications for policy design from it. This body of work builds on brand new data sets, such as modern-day data (e.g., European patent office data since 1975) or historical long-run data (e.g., the universe of all U.S. inventors since 1836). It also leverages new theoretical and structural methods and models that build up behaviors from the micro-level of the firm all the way to their macro growth implications.

Section 3 considers how the quantity and quality of innovation respond to tax policy; Section 4 focuses on the geographic mobility of innovation and inventors across U.S. states and countries. Section 5 focuses on the declining business dynamism in the U.S. and how specific policies can improve firm entry and productivity. Section 6 studies the effects of tax policy on the quality composition of firms, inventors, and teams and how the right design of policy can allow policy makers to foster the most productive firms without wasting public funds on less productive ones. Section 7 shows how policy can orient research into different directions, e.g., from applied to basic research, or from dirty technologies to clean ones.

¹When empirical work is presented, the methods are described in some detail in order to allow the reader to better assess the reliability of the estimates.

2 Through What Channels Do Tax Policies Shape Innovation?

In this section, we conceptually map the effects of different tax policies on innovation, emphasizing the many channels through which policies can play a role. Each of the channels in this figure will be discussed in light of the existing literature below. To organize thoughts, Figure 1 gives a one-glance schematic representation of the framework.

The main actors for innovation. Innovation is done by firms or individual inventors. These key agents of innovation are represented at the center of the column. Inventors can be self-employed or work in companies' R&D labs.

Key characteristics of firms and inventors that have to be considered by policy are represented on the diagram: Inventors and firms can be of varying productivity, i.e., the efficiency with which they convert R&D and research inputs into innovations. The productivity composition of firms and inventors will shape the impacts of various policies, as well as endogenously be affected by them. For firms, it is not just the quality of the firm overall that matters, but also more specifically as emphasized by the literature below, the quality and composition of its research teams. Firms can also be at different stages in their life cycle, from early start-ups to mature, large firms. Similarly, inventors can start off as young, inexperienced inventors, and improve their skills through learning and experience over time.

Innovation inputs and actions. The production of quality and quantity of innovation requires inputs. These inputs are either tangible – lab space, equipment, or material resources – or intangible – effort, the skill and know-how of workers, or the efficiency of management.

As represented by the bubbles in the figure, inventors and firms each have many possible margins on which to optimize, all of which could in principle be responsive to many different policies. Inventors first make an occupational choice: whether to become inventors at all. They also need to decide whether they want to be self-employed or employed by companies. They also choose where to locate geographically. Whether they work for companies or not, they have to select their level of tangible and intangible inputs. Once an innovation is created, they have to choose whether to sell it to a firm or rather to incorporate and build a business around the new invention. The innovation and the associated flows of income can thus move from the personal to the corporate sector.

Companies have to choose whether to enter a given market, remain in operation, or exit. They also select their geographic location, which could be in multiple places and different for production or research facilities. Companies decide on their innovation and R&D inputs and their research employment. They also choose whether to direct their research operations towards basic or applied innovation (?). Basic innovation is defined as the “systematic study to gain more comprehensive knowledge or understanding of the subject under study without specific applications in mind,” while applied research is a “systematic study to gain knowledge or understanding to meet a specific, recognized need.” In addition, they have to decide whether to engage in internal innovations – defined as improvements to existing products – or external innovations – defined as the creation of new products or overtaking of competitors' products.

General and Specific Tax Policies. The main tax policies that shape innovation can be classified into general tax policies, such as personal income taxes, corporate income taxes, or education subsidies, and more targeted innovation-specific tax policies, such as R&D tax credits, start-up subsidies, research subsidies for specific types of research and R&D, and location-specific incentives for firms and inventors.

Regarding general tax policy, inventors and firms could in principle be affected by personal and corporate income taxes. For inventors, the personal income tax directly affects the size of their post-tax income. For self-employed inventors, the corporate income tax matters if and when they decide whether to incorporate or not. For employed inventors and firms, any surplus-sharing implies that both the personal and corporate tax shape the payoffs. The extent to which the corporate income tax will affect firms' R&D decisions depends on the share of research inputs that can be expensed; with full expensing, corporate income taxes should have no effect on R&D investment decisions. It also depends on the presence and size of fixed costs, which have to be recouped through future net-of-tax income flows.² Firms will also take into account the personal income tax when deciding how many researchers to employ if they have to pay some compensating differential when and where taxes are higher. Education subsidies can shape the choice to acquire the skills needed to become a high-skilled inventor. Corporate and non-corporate inventors can have different responses to tax policy, both because their payoffs may be differently affected by it and because they may have different motives to engage in innovation (which may be why they are in the corporate sector or not to start).

Turning to the more targeted policies, R&D tax credits can affect the full range of decisions made by firms, and can change the relative payoff to inventors from incorporating and working for companies. Targeted subsidies to start-ups can favor entry; subsidies for specific types of research such as applied or basic can affect the direction of innovation; and location-specific policies can attract firms and inventors to certain places.

Responses to Tax Policies. The elasticities to all these different tax and subsidy policies will be a composite of the behavioral elasticities and the technological elasticities. Behavioral elasticities measure how firms and inventors adjust all their margins of action; technological elasticities capture how sensitive innovation outputs—quantity and quality—are to each of these actions. On the technological elasticities, one can imagine two polar extremes: at one end is the case of “Newton sitting under the tree;” the apple falls, innovation happens entirely inelastically. At the other end would be a very mechanical innovation process where more inputs would automatically translate into more output, e.g., if testing many more new chemical combinations results in a scaled up probability to find a new material. Similarly, on the behavioral elasticities, one can imagine the polar extremes of the “mad genius” who is only doing innovation for the love of science and the purely profit-driven entrepreneur. Thus, the elasticities of all innovation actions, and, hence, of the resulting innovation, are empirical questions.

Dynamics. Innovation is an investment-type activity that involves forward-looking behavior, as upfront costs today potentially yield a stream of benefits in the future. Thus, inventors and firms need to form some expectations about the net present value of those returns, which are shaped by the range of aforementioned policies and their predicted changes over time. If firms expect corporate tax rates to increase in the near-future, for instance, the net present value of their payoff from innovation would be reduced relative to a scenario in which they expect the corporate tax to decrease. Tax policies are hard to predict and thus expectation formation becomes a key issue for agents deciding whether and how much to engage in innovation.

In addition to these forward-looking effects, there can also be lags with which innovation responds to changes in policies. Innovations take time to produce and the time span between the

²Given the empirical evidence below, it is likely that there is less than full expensing and/or fixed costs, as the corporate income tax does matter. It is likely that many R&D inputs are either unobservable (such as the intangible inputs) or hard to measure.

changes in behavior and the creation of an innovation is also an empirical question. The lags could be different for different policies. For instance, as we will show below, education policies take much longer to work than do R&D tax credits; carbon taxes work faster but at a higher cost than do research subsidies for clean innovation.

Micro to macro: Individual and economy-wide responses The responses to tax policy measured depend on the level of analysis. At the individual firm and inventor level, all the response margins just described could in principle respond. Zooming out to the more macro level, e.g., localities or states in the U.S. or countries in the world, additional effects will be layered on the micro-level effects. For instance, factors can shift between places, leading to a reallocation that may or may not add value overall. To take the example of U.S. states, part of the response seen at the macro state-level could be due to pure business-stealing and cross-state spillovers, without a corresponding increase in innovation at the federal level. The same holds for international movements of factors. In addition, at the macro level, the effects of taxes can be augmented or dampened by other features related to tax policy, such as the research infrastructure or amenities, and the formation of innovation hubs. When reading the empirical literature, it is important to bear in mind the level of analysis and avoid extrapolating without care to other levels.

Measuring innovation: How can we best measure innovation and growth? The literature described in this chapter uses mostly patent data – from the European patent office (EPO), the U.S. patent office (USPTO), the international Patent Cooperation Treaty, or from historical patent records – to measure the quantity of innovation. However, some of them instead use firm growth or R&D spending to proxy for innovation. It is of course impossible to measure “all of” innovation systematically with any of these measures. Take patents for instance: a share of inventions is not patented. By their very nature, patents will be highly correlated with the quantity of innovation and are thus a prime measure that can shed a lot of light on the issues surrounding innovation, sometimes in conjunction with the other aforementioned measures.³ To measure the quality of innovation, an often-used metric is the forward-citations going to a patent, which have been shown to be a proxy for economic value and instructive about the importance of the invention for subsequent innovation (??). The length of patent claims are also used to measure whether an innovation is incremental or radical (?).

3 Quantity and Quality of Innovation

Turning to the literature, we start by outlining some recent findings on the effects of general taxation (personal and corporate income taxes) on the quantity and quality of innovation.

The Study. The United States has experienced major changes in its tax code throughout the twentieth century. Have these tax changes influenced innovation at either the individual or corporate level? This challenging question has largely gone unanswered because of a lack of long-run systematic data on innovation in the U.S. and the difficulty in identifying the effects of taxes. ?

³? find that the share of innovations that are patented is very low for low-tech industries such as textiles (8.1%) where it is mostly process innovation and high for high-tech industries such as pharmaceuticals (79.2%). Petra Moser believes that historically, the share patented has been around 50% (<https://www.knowablemagazine.org/article/society/2018/do-patents-invent-innovation>). In this Brookings paper volume summary from 1989, Mansfield reports findings that 60% of innovations in the auto industry and 80% in the pharmaceutical industry are patented. (https://www.brookings.edu/wp-content/uploads/1989/01/1989_bpeamicro_summary.pdf).

construct and leverage three brand new datasets from historical data sources to shed light on these questions. These new data sources are i) a panel of the universe of U.S. inventors since 1920 and their associated patents, citations, and firms; ii) a panel of all R&D labs in the U.S. since 1921, matched to their patents and with data on their research employment levels and location; iii) and a historical state-level corporate and personal income tax database. This unique combination of data allows them to systematically study the effects of both personal and corporate income taxation since 1920 on individual inventors and individual firms that do R&D – i.e., the “micro” level– and on innovation at the “macro” state-level.

The innovation outcomes are, for instance, the quantity of innovation, as captured by the number of patents, the quality of innovation (as measured by patent citations), or the share of patents assigned to companies rather than individuals at both the macro state level and micro individual inventor level.

It is challenging to convincingly identify the effects of taxes on the quantity and quality of innovation because when general tax policy changes in a state it may be in response to changes in economic conditions and it may occur contemporaneously to other policy changes, both of which could also affect innovation outcomes independently. Therefore, the authors approach the question from several angles, which all yield consistent results. First, they control for a detailed set of fixed effects, including state, year, and at the individual level, inventor fixed effects, plus individual or state level time varying controls; these go a long way towards absorbing unobserved things that vary by state, by year, or by inventor. In addition, they exploit within-state-year tax differentials between people in different tax brackets (e.g., the top tax bracket versus the median one). This allows them to control for things that vary at the state and year level and to filter out other policy variations or economic circumstances that may occur at the same time in the state. Second, at both the macro and micro level, they use an instrumental variable strategy that consists in predicting the total tax burden facing a firm or inventor – which is a composite of state and federal taxes – with the changes in the federal tax rate only, holding the state taxes fixed at some past level. This provides variation that is only driven by federal level changes, and thus, exogenous to any individual state. Third, they compare the evolution in innovation in neighboring counties that lie in different states (so-called “border counties”) that are economically and geographically similar, but face different state-level taxes.

Key Findings. The paper finds that both corporate and personal income taxes negatively affect the quantity and quality of innovation at the macro state level and the individual inventor and firm levels. The elasticities to taxation of all these innovation outcomes are relatively large at the micro level, of orders of magnitude similar to those of taxable income elasticities estimated. For instance, at the individual inventor level, the elasticity of patents to the personal income tax is 0.6-0.7, and the elasticity of citations is 0.8-0.9. This means that a 10 percent change in personal income taxes will reduce patents produced by 6 percent and the quality to a similar extent.

When it comes to the dynamic effects, innovation responds to general personal and corporate income taxes with a lag: the response starts one year after the tax change and increases for the next three years. Although, as described above, there could also be forward-looking effects, since innovation is an investment-type activity that will potentially yield a stream of returns for a period in the future, no “lead” effects of the tax rates are observed in the data. This could be because, on average, current tax rates may be the best predictor of future tax rates.

Zooming out from the individual inventor level and to the macro “state level,” responses to taxes are close to double that magnitude. They are partially – but not entirely– caused by cross-state spillovers and business stealing, which does not add to total innovation in the U.S., but simply

reshuffles resources among states.

The paper also shows that corporate inventors, and, consequently, corporate patents, are consistently more elastic to personal and corporate income taxation. This could be because corporate inventors have payoffs which are more sensitive to taxes and because they have different motives for innovation. Finally, and crucially, agglomeration effects appear to matter significantly. Inventors are less sensitive to taxation where there is already more innovation in their field in a given location.

Policy implications: The quantity and quality of innovation appear to have been responsive to taxation throughout the 20th century. In magnitudes, the responses at the individual firm and inventor levels are somewhat larger than other standard margins we typically take into account when setting tax policy, such as labor supply or the overall taxable income elasticity. This means that the efficiency costs of general taxation in terms of innovation should be taken into account in tax evaluation. In addition, policy makers and analysts need to be very careful in extrapolating from state-level responses to federal-level responses. As emphasized, state-level responses are inflated by cross-state spillovers which are zero-sum effects from the federal point of view. The better approach is to start from the micro-level elasticities and aggregate these up to the federal level.⁴

4 Mobility of Inventors and Firms

Another margin along which both firms and inventors can respond to tax policies is their location choice. Recent evidence shows that there is tax-induced mobility both across U.S. states and across the major patenting countries in the world.

Historical mobility in the U.S. since 1920. The historical project described in Section 3 shows that inventors move in response to state personal income tax rates. The elasticity to the net-of-tax rate of the number of inventors residing in a state is 0.11 for inventors who are from that state and 1.23 for inventors not from that state. Inventors who work for companies are particularly elastic to taxes. Crucially, agglomeration effects appear to matter for location as well: inventors are less sensitive to taxation in a potential destination state when there is already more innovation in that state in their particular field of inventive activity.

International mobility since 1975. There is a heated public debate about whether higher top tax rates in a country will cause a “brain drain” of high income and high-skill economic agents. In fact, many of the great inventors were international immigrants: Alexander Bell, inventor of the telephone and founder of the Bell Telephone Company; James Kraft, inventor of a pasteurization technique and founder of Kraft Foods Inc.; Ralph Baer, creator of a TV gaming unit that launched the video game industry.

Inventors are much more mobile than other high-skilled individuals. Thus, they carry and transmit their valuable knowledge and expertise to others (as shown in Section 6.4), making them essential not only for new knowledge creation, but also for its diffusion. Yet, until recently, little was known about the international mobility of labor in response to taxation and rigorous evidence was lacking because of a scarcity of international panel data. ? close this gap by using a unique type of international panel data on inventors from the European and U.S. Patent offices, as well as from the Patent Cooperation Treaty to study the international migration responses of superstar inventors to top income tax rates for the period of 1977-2003. They are able to tackle one major challenge

⁴This will require to be rigorous a structural model of how individual level responses of firms and inventors map into federal-level outcomes and the incorporation of general equilibrium effects.

that arises when studying migration responses to taxes, namely, to model the counterfactual payoff that an inventor would get in each potential location.⁵

Their identification of the effects of the top tax rate relies on filtering out all country-year level variation and exploiting the differential impacts of the top tax rate on inventors at different points in the income distribution within a country-year cell. To implement this strategy, superstar inventors are defined as those in the top 1% of the quality distribution, and similarly construct the top 1-5%, the top 5-10%, and subsequent quality brackets. It is known that inventor quality is strongly correlated with income and that top 1% inventors are very high up in the top tax bracket. The probability of being in the top bracket and the fraction of an inventor's income in the top bracket declines as one moves down the quality distribution. Top 1% inventors and those of somewhat lower quality are comparable enough to be similarly affected by country-year level policies and economic developments; but only those inventors in the top bracket are directly affected by top taxes. Hence, the lower quality, top 5-10%, top 10-25%, and below top 25% groups serve as control groups for the top 1% group.

The paper finds that superstar inventors' location choices are significantly affected by top tax rates. The elasticity to the net-of-tax rate of the number of domestic superstar inventors is around 0.03, while that of foreign superstar inventors is around 1. These elasticities are larger for inventors who work for multinational companies. On the contrary, inventors are less sensitive to taxes in a country if their company performs a higher share of its research there, suggesting that the location decision is influenced by the company and by career concerns that may dampen the effects of taxes.

Policy implications: Tax-induced mobility is an issue to take into account, especially when it comes to highly-mobile and highly-skilled people such as inventors who can make major economic contributions to their country of residence. But the right answer may not be to slash general taxes and to engage in relentless tax competition, e.g., through preferential tax regimes for foreigners as has been done in the E.U. As argued in ?, these are "beggar-thy-neighbor" policies which can reduce overall welfare; international or cross-state tax cooperation is much more fruitful in the long-run.

So what can be done then? One important margin along which states and countries can act is to provide better amenities and infrastructure for innovation. The studies above consistently show that agglomeration effects significantly dampen the responsiveness to taxes. This occurs for two reasons that can be leveraged. First, agglomeration effects are a proxy for the quality of the research amenities and infrastructure in a place, which are valued by many inventors who choose to live there. Second, inventors and firms directly benefit from being around like-minded talented innovators. This implies that attracting innovation to a locality in the first place can start a virtuous circle that continues into the long-run. Thus, one way of being able to continue using general tax policy for its intended purposes of raising revenues and redistributing income without stifling innovation and causing the outflow of talent would be through the better provision of amenities and infrastructure for innovation.⁶

5 Declining Business Dynamism in the U.S.

As highlighted in the introduction, business dynamism in the United States has been slowing in the last several decades. The key "ten facts" highlighted in the literature (?) are that market

⁵This is thanks to a set of detailed controls that come from the patent data, notably measures of an inventor's quality, based on his past citations.

⁶In fact, the tax revenues themselves go towards the investment in such amenities.

concentration, average profits, and markups have risen; the labor share has decreased; the labor productivity gap between frontier and laggard firms has increased; entry rates have declined, as has the share of young firms; job reallocation has slowed down; the dispersion of firm growth has increased and the rise in market concentration and the fall in labor share are positively correlated.

The decline in business dynamism is intrinsically linked to the life cycle of firms. It is thus critical to understand the distinct innovation strategies that firms follow at different stages of their life cycles. Young startups and small firms initially explore radical new ideas; for those that survive and grow bigger, innovations tend to become more incremental. Both firms and individuals can also be tempted to start erecting barriers for their competitors and prevent entry into their market. They can do this through political connections and lobbying as they gain power, or directly through the intellectual property rights (IPR) system. This can slow down innovation and influence business dynamism.

Policies can affect firms' innovation very differently based on where they stand in their life cycle. For instance, existing research tax credits help mostly big and profitable firms, as tax credits are only beneficial for firms that make positive profits. Many countries have special policies for start-ups and others have targeted policies for small businesses. In this section, we consider the effects of various tax policies based on different aspects of firms' lifecycles and their implications for business dynamism in the U.S.

5.1 Start-ups and Venture Capital

Venture capitalists play a major role in the screening, monitoring, and financing of startups. ? show that VC-backed startups grow faster in their early stages, produce innovations of better quality, and have a higher likelihood of becoming large firms and produce high quality innovations over time. These effects are stronger if the firms are matched to more experienced venture capitalists. The authors estimate that the existence of venture capital and an efficient match between startups and VCs is important for innovation and growth.

When it comes to tax policy, empirically, VC-funded firms are de facto taxed at preferential rates relative to non-VC-funded ones. The authors show that increasing the tax rate on VC-funded firms to harmonize it with the tax rate on non-VC-funded ones would have a significant negative impact on aggregate innovation. Because VCs have a significant value-added to the startup and innovation process and do not simply crowd-out other firms, the reduction in their activity that would be caused by a higher tax rate is not offset by an equivalent rise in innovation and success of the non-VC-funded firms.

Policy Implications: VCs contribute significantly to the innovation by fostering high-quality startups. Maintaining a low tax burden on VC-funded firms could possibly foster innovation.⁷

5.2 Small versus Large Firms

Firms can adopt different innovation strategies. They can engage in external innovation, which creates new products and captures market share from competitors and they can engage in internal innovation which improves product quality for the products lines that the firm currently owns. Small and large firms typically choose quite different routes when it comes to innovation.

⁷Naturally, in a world without government budget constraints, all firms should be taxed at low rates. The statement here is that given a scarcity of government funds and if the goal is to foster innovation, tax cuts should be given preferentially to VC-funded firms rather than to non-VC-funded firms.

? explore this key distinction and provide an explanation for why the data shows that small firms experience faster growth on average and contribute disproportionately to big and radical innovations. This is not due to different capabilities of firms by assumption, but is rather the result of structural estimation of the model. The authors quantify their model using U.S. data from the Census Bureau for 1982-1997, finding that decreasing returns to external innovation in larger firms to be an important departure from the perfect scaling of the ? framework. This work allows the authors to conclude that growth impacts of external innovation have exceeded internal innovation for the recent U.S. economy, which in turn helps identify some of the special role that small, innovative firms and new entrants can play in economic growth.

When thinking of total innovation in this context, there are three sources: external innovation by incumbents, internal innovation by incumbents, and (by necessity) external innovation by entrants. While all of them matter, the authors estimate that the key channel for growth is external innovation by incumbent firms. This innovation is done disproportionately by small firms.

Policy implications: Small firms find it more beneficial to engage in external research which leads them to produce disproportionately many more radical, important innovations. Preferential tax rates and tax breaks for small firms can foster this phenomenon further and improve the quality of innovation and prevalence of breakthrough innovations. In fact, for different reasons, many countries have some sort of preferential tax treatment for small enterprises.

5.3 Political Economy

As firms establish themselves in the labor and product market, they may start being tempted to keep competitors out. One way to do this is through political connections and lobbying. ? show that this happens very frequently in Italy and that it reduces the dynamism in the economy, leading to less reallocation of resources, less innovation, and less entry. Although their empirical analysis uses Italian data, it carries lessons for the U.S. as well.

To some extent, political connections can serve a productive role if they alleviate some bureaucratic and regulatory burdens that cause inefficiencies. But they are also costly and require firms to expand resources to maintain them. Larger firms will find it more profitable to incur these costs and to be politically connected. The authors show that there is a leadership paradox: the leading firms in each industry (i.e., those with the largest market share) absorb the most resources, but are also the least innovative ones, relative to their size. Firms that become connected temporarily enjoy higher employment and sales growth, but lower productivity growth. Over time, industries in which the incumbents are politically connected become more sluggish, with less entry as new entrants are discouraged to try and compete not only in terms of productivity, but also in terms of the regulatory and bureaucratic burdens. Politically connected incumbents thus have an advantage. Since political connections by incumbents discourage entry, incumbents may decide to preemptively become politically connected to shelter themselves from competition. Politically connected industries will be dominated by older and larger firms and feature low innovation and productivity growth.

Policy implications: If political connections are predominant and cannot be directly prevented by policymakers, tax policy could take on a corrective role and compensate for the disadvantage that small new entrants face relative to the large, politically connected incumbents. If the frictions are caused by taxes to start with, they could be removed directly, thus also removing the incentive to become politically connected. A more indirect way would be to tax larger incumbents at higher rates to give new entrants a chance to compete.

5.4 Declining Business Dynamism and Innovation

? provide a theoretical and quantitative model that is able to explain the ten facts listed above and offers an explanation for what has happened to business dynamism in the U.S. The dominant force driving these patterns is a decline in the rate of knowledge diffusion from the frontier, most-advanced firms to the laggard ones. This force can explain all ten facts in a way that other explanations cannot. In addition, there is direct evidence for it. For instance, patenting has become more concentrated among firms with many patents. The nature of patents has also changed since the 2000s, with longer claims (indicative of more incremental, rather than radical innovations) and more self-citations. Overall, the evidence is consistent with a use of intellectual property protection by leading firms to limit knowledge diffusion and entrench their market power.

In their model, in each sector, two firms compete for market leadership. One firm represents “the best”; the other, “the rest.” Prices and markups are a function of the technology gap between firms, as the leading firm can charge up to the level where the non-leading firm with the worse technology can no longer capture a sizable market share. The incentive for leaders to keep innovating is to grow the gap in technology between them and their followers and be able to increase prices and markups. Existing followers’ incentives to innovate are to catch up with and leapfrog the leader to capture more of the market; similarly, new entrants’ incentive is to one day become the market leader.

A key feature of this model, which is clear in the data as well, is that firms make strategic innovation investment decisions, taking into account where they stand relative to others. When firms are neck-and-neck and compete very intensely for market leadership, there will be a lot of innovation and business dynamism. But when leaders do very well and open up the technology gap between themselves and their followers, prospects for entrants and followers become dim. Then, innovation efforts are reduced and entry declines. Because of this, when the diffusion of knowledge declines, market leaders are protected and establish strong market power. This in turn discourages followers and entrants and slows innovation in more concentrated sectors. As entry and the threat from competition are diminished, market leaders also slow down on their innovation efforts. Overall, business dynamism and innovation decline.

Corporate taxes have played a role in this evolution too, but it is only a minor one. Corporate tax increases can explain around 10 percent of the decline in business dynamism. This is because lower taxes, which increase net-of-tax profits and will thus only matter for firms that have some market share. They will hence only affect followers and potential entrants in a very muted way, as they are far from taking over the leader and have to discount the potential future gains heavily.

Policy implications: The key lesson from this recent set of papers is that powerful, large incumbents are using their market power in order to prevent entry and competition by followers. The most direct policy approach to this would be through regulation and competition policy. However, corporate tax policy can play a (second-best) role too if it can be designed so as to impose a heavier burden on larger, more mature companies rather than on new entrants. As shown above, this will only have a very small disincentive effect on new entrants (facing the prospect of one day becoming market leaders) but could eat away at part of the advantage the incumbent market leader enjoys.

6 The Composition and Quality of Inventors, Firms, and Teams

Firms and inventors are not all equally effective in producing innovations. The quantity and quality of innovation in an economy will depend on the composition of firms and inventors, which can also

be affected by general and targeted tax policies. In this section, we explain how different policies can play on that margin, considering in turn firms, inventors, and teams within firms.

6.1 Firms: Reallocation of Resources between Firms

Not all firms are equally efficient at producing innovation. Whether because of the quality of their ideas, management, or workforce, some firms are excellent at turning research inputs into major innovations, while others are not. Tax policies can affect the selection of firms, their entry and exit, and the reallocation of resources across good and bad firms. ? build and estimate such a dynamic firm-level innovation model using U.S. Census micro firm-level data and patent data. A key finding is that taxing incumbent firms can be very beneficial, in addition to subsidizing R&D. R&D is subsidized to correct for the underinvestment in innovation due to non-internalized spillovers. In the presence of fixed costs, taxing the operations of incumbents encourages less productive firms which are closer to the exit margin to leave the market. This frees up valuable resources – i.e., skilled researchers – for more productive firms to hire. On the other hand, when it is not possible to have type-specific R&D subsidies targeted to good versus bad firms, uniform subsidies to R&D alone will not achieve this positive selection because they will benefit both low and high productivity firms and encourage low productivity firms to survive, grow, and absorb scarce resources.

Policy implications. In addition to raising revenues, corporate taxes can serve an efficiency purpose too. A sizable, uniform tax on incumbents combined with a uniform R&D subsidy can improve growth and increase welfare. This is an important finding that shows that, despite the distortionary effects of taxes, they can serve an allocative role, by selecting firms that are good enough to survive despite taxes. Corporate income taxes can have a cleansing effect on the economy, freeing up valuable innovation resources for the most productive firms.

6.2 Firms: Optimal R&D Policies

Policy makers can do even better with non-uniform policies, by explicitly trying to screen good firms from bad ones. The major obstacle to doing so efficiently is asymmetric information – a key feature in the innovation arena. The innovation literature has extensively addressed how to deal with spillovers, but it has much less dealt with asymmetric information about firms and how to separate productive from nonproductive firms. Yet, as a large empirical literature shows, the quality of a firm’s organization, management, processes, or ideas – which shape its innovation outcomes, conditional on inputs – are private information and very difficult for outside parties, including the government, to observe. The literature documents the manifestations of asymmetric information in innovation, between firms and their shareholders or investors; this problem is even more pronounced between firms and the government. In addition, in the patent and firm data, if one tries to predict the innovation quality of a firm, the prediction is very poor. Seeing which firms are good at innovation and which ones are not is inherently difficult, even using a very large set of observables that is likely a very generous overestimate of what the government could realistically condition policies on.

One way to address the asymmetric information problem is that adopted by Venture Capitalist firms, which perform very hands-on and thorough screening, and provide staged financing subject to intense monitoring. But this intensive hands-on approach is not easily scalable and thus not applicable when it comes to thinking about large-scale government policies. Instead, what the government can do is to set decentralized tax and subsidy policies that can vary non-linearly with

profits and R&D investments, and in such a way that firms of different productivities will pick their customized, uniquely efficient levels of investments and production.

? tackle this problem using a new dynamic mechanism design approach. The key feature of their analysis – and the main impediment to fixing the market distortions typical to innovation such as nonappropriability and spillovers in a non-distortionary way – is that firms are heterogeneous in their research productivity, and importantly, this research productivity is private information and unobservable to the government. A higher research productivity allows a firm to convert a given set of research inputs into a better innovation output. In addition, while some of the inputs into the R&D process are observable (so-called “R&D investment”), others are unobservable (“R&D effort”). The firm’s research productivity also evolves stochastically over time. Although the firm has some advance information about its future productivity, it cannot perfectly foresee it. As a result, at the time when the firm invests resources in R&D, the innovation outcome that will result from these R&D investments is yet uncertain.

The authors’ main findings are as follows. Asymmetric information can significantly change the optimal policies. From a theoretical perspective, the constrained efficient incentives for R&D trade-off a Pigouvian correction for the technology spillover and a correction for the monopoly distortion against the need to screen good firms from bad ones. How much R&D should optimally be subsidized depends critically on a key parameter, namely the complementarity of R&D investment to R&D effort (i.e., the complementarity between observable and unobservable innovation inputs) relative to the complementarity of R&D investment to firm research productivity. The more complementary R&D investment is to firm research productivity, the more rents a firm can extract if R&D investment is subsidized. This puts a brake on how well the government can set the Pigouvian correction and correct for the monopoly distortion. Optimal screening in this case requires dampening the first-best corrective policies. On the other hand, if R&D investments are more complementary to unobservable firm R&D effort, they stimulate the firm to put in more of the unobservable input, which is unambiguously good and would make R&D subsidies optimally larger. Other key determinants of the magnitudes and age-patterns of the optimal policies are the persistence of firms’ research productivity shocks and the strength of spillovers.

The data shows that R&D investments are highly complementary to firm research productivity: highly productive firms are disproportionately good at transforming R&D inputs into innovation. Given that this implies that higher productivity firms have a comparative advantage at innovation, it is better to incentivize R&D investments less for the lower productivity firms, as this makes mimicking them more attractive for high productivity firms.

Policy Implications: It is possible to very closely implement the optimal allocations with simple non-linear or linear policies that feature lower marginal corporate income taxes for more profitable firms and lower marginal subsidies at higher R&D investment levels. The policies can be simplified even further without much loss, as the most important quantitative feature turns out to be the nonlinearity in the R&D subsidy. Thus, making the profit tax linear only generates a small welfare loss. The intuition is that a constant profit tax that is set at a too generous level for low profit firms and at about the right level for high profit firms does reasonably well, since the loss from giving low profit firms a too generous tax is quantitatively small given that they make low profits to start with. Therefore, linear corporate income taxes such as the ones we see around the world can be very close to optimal for innovating firms if combined with the right nonlinear R&D subsidy.

6.3 Inventors and Education Policies

When it comes to the composition and quality of individual inventors, innovation policies have an important interaction effect with education policy through the occupational choices of inventors, and hence, the supply of high-skilled researchers. ? point out that education policy and general or targeted tax policies for innovation will address different frictions in the innovation chain. In their setting, inventors of varying abilities and with different career preferences take time to build their human capital and face financial constraints in acquiring education. As a result, in the short-run, targeted policies such as R&D policies may not be as effective as expected; they may face the bottleneck of insufficient supply of research talent due to lack of education capacity or credit constraints. In the long-run, these policies may still have limited effectiveness if they are not coupled with education policies. This new interplay can explain why innovation models typically predict much larger effects of R&D policies than those that are observed in the data.

Policy implications: How do different policies influence the aggregate innovation and economic growth? The authors find that the impact of R&D subsidies can be strengthened when combined with higher education policy that sorts talented but credit-constrained individuals into research. In addition, the role of education policy is increased in societies or times when financial constraints on the acquisition of education are more stringent. Education subsidies are particularly critical and effective in unequal societies where many individuals face financial constraints preventing them from efficiently acquiring education. In these cases, R&D policies alone are highly ineffective.

There is, of course, a critical timing issue. In the short-run, only R&D policies can be effective, as education policies act with longer lags. R&D policies stimulate the purchase of more research capital and equipment, making researchers more productive almost upon impact. However, educational slots expansion takes some time and surpasses R&D after six years. Educational subsidies, on the other hand, take the longest to transmit to the growth rate, but gradually become the most effective policy tool in the long-run

6.4 Teams and Knowledge Diffusion

Inventors do not work alone: most patents are the result of collaborative work and are produced by teams of inventors of different talents and skill-levels. In addition, inventors learn from each to produce better innovations. When an inventor interacts with other inventors who are more knowledgeable than them, they also improve their knowledge and produce subsequent higher quality innovations.

? provide a model and empirical analysis that captures these key features observed in their data. In their framework, inventors can learn (i.e., improve their productivity) in two ways. They can meet others and interact with them and they can learn on their own, through learning-by-doing, formal education, experience, or individual discovery. Given their realized productivity after learning, inventors form teams. Some inventors who are highly-productive and knowledgeable become “team leaders” and work with less-skilled team members to produce innovations. Better team leaders will be able to hire larger teams and produce better innovations. The quality of innovation, and thus, technological progress, will depend on the quality of teams in the economy. The authors estimate the model using new European Patent Office data for inventors across many years and countries and find that interactions with others are quantitatively very important for improving inventors’ productivity, and hence, economic growth. Interactions can occur at the firm-level, technological field-level (in a given area), or at different geographic levels. In addition, there is a strong complementarity between access to external knowledge and learning from others: if others

around oneself learn more from outside sources and then interact more, one will end up interacting with more knowledgeable people and learning more as well.

Therefore, when thinking of the effects of tax policies, one has to consider, on the one hand, their implications for team formation and team composition, and on the other hand, their implications for interactions, learning, and knowledge diffusion among inventors.

Policy implications: In line with the mobility results discussed above, policies that attract many inventors to a given area can foster interactions and thus learning. As discussed above, preferential local tax regimes can achieve this goal, but at the cost of penalizing other areas in a zero-sum way. Better amenities can achieve the goal more efficiently without ruthless tax competition. On the other hand, large employer payroll taxes or firing taxes, which are particularly heavy in many European countries, can reduce labor market fluidity and prevent inventors from moving to their best matching team. Education subsidies increase the quality of the pool of inventors and make it more worthwhile for inventors to try and interact with others to learn from them.

7 Applied versus Basic Innovation and Choice of Technologies

Innovations come in different shapes and sizes. Companies and inventors can choose to orient their research in different directions and their choices here too can be shaped by tax policy.

7.1 Applied and Basic Research

A major distinction, as described in Section 2, is between basic and applied research. ? refer to “Pasteur’s quadrant” for illustration of the different types of research. At one extreme, lies pure basic research, as done most often by the public sector in academic institutions and universities. At the other extreme lies purely applied research, destined for immediate commercial use. In between exists a mix of basic and applied, captured in their paper by the private sector’s basic research – research that is ultimately driven by a profit motive and with the hope of being one day applicable, but without immediate, intentional commercial implications.

Without government intervention, there is a stark misallocation of research effort. The authors find that 68% of the spillovers from basic research are not internalized. Once these different types of research are taken into account, it appears that the bigger problem is not the insufficient investment in research overall, but rather the wrong allocation of research efforts between basic and applied innovation. There is too little investment in basic research – yet, there can be too much investment in applied research in the face of competition between firms, if there is strategic complementarity between basic and applied research (i.e., the returns to applied research increase in the level of basic research).

The authors also shed light on the debate about the worrying decline in research productivity in the U.S. economy. They highlight the strong complementarity between public and private research efforts. When there is more investment by public entities in predominantly basic research, private research investment becomes more productive too and increases.

Policy Implications: Which policies can address the inefficient allocation of research efforts? A uniform research subsidy to private firms – that subsidizes all of their research at the same rate – carries a large fiscal cost in this context. Although it will stimulate investment in basic research, it will generate even further over-investment in applied research. Subsidizing applied and basic research at different rates can reduce fiscal costs without compromising investment in innovation.

The authors find that the optimal subsidy rate for basic research in their model is almost five times as high as that for applied research. Clearly, distinguishing between applied and basic research inputs could be difficult, which means that it's important to allow for some misclassification by firms, which will be tempted to relabel applied research as basic research. But even with substantial levels of misreporting of research types, a higher subsidy for basic research remains very effective. Going forward, finding a feasible way to differentiate between basic and applied research is essential to better innovation tax policies.

In addition, subsidies and funding for public research could also indirectly foster investment in private research, as public research is highly complementary to private research.

7.2 Innovation in Green Technologies

Tax policies can also orient research into different directions when it comes to the environment and the development of clean technologies. Given how pressing and critical of an issue climate change is, tax policy tools for innovation in these areas have to be very carefully considered and deployed.

? provide a clear theoretical and quantitative framework to think about tax policy for innovation in clean technologies. Goods can be produced with either a “dirty” (polluting) technology or a “clean” (less polluting, environmentally-friendly) technology. Producers choose which technology to employ based on their costs which in turn depend on how efficient these technologies are and based on the public policies, such as production taxes that vary by technology type. For instance, carbon taxes or taxes, on other polluting particles or greenhouse-gases would imply a higher tax on the dirty technology. In addition to picking their production technology, private firms can also choose to do research to improve either the clean or dirty technology. Research and innovation decisions are also shaped by public tax policies and the current state of technology. If the clean technology is very far behind the dirty technology in terms of efficiency, research in it that produces only incremental improvements is unlikely to yield immediate benefits to producers and is thus not profitable in the short-run. However, sustained research efforts and cumulative incremental improvements may eventually render the clean technology competitive and profitable.

Taxing pollutants, e.g., through carbon taxes, can redirect research towards the clean technologies. Subsidizing clean research can also achieve this goal. But as long as the dirty technology remains much cheaper than the clean one, carbon taxes will reduce pollution but at a high efficiency cost in terms of foregone consumption. A research subsidy can redirect research successfully, even when combined with a low carbon tax initially, until the clean technology becomes able to compete with the dirty one. Research subsidies in this model can be optimal even if there is no underinvestment in research overall; they are used to counteract the negative externalities of pollution for the environment.

Policy implications: Both carbon taxes (as well as taxes on other pollutants) and research subsidies for clean technology can be used to direct innovation to clean technologies. However, carbon taxes are very costly initially when the clean technology is still very inefficient relative to the dirty technology. Thus, initially carbon taxes are a more cost-effective tool to correct for the direct externality of pollution from carbon or other pollutants, but it is research subsidies that are more cost-effective for guiding research towards clean technologies. A mix of policies that stimulates investment in green technologies at lower fiscal cost than other policies can be described as follows: Policy initially focuses heavily on the research subsidy, which declines over time, Carbon taxes are backloaded initially (increase over time as the clean technology becomes more efficient), but eventually also decline as pollution is reduced thanks to the use of cleaner technology.

8 Conclusion

Tax policies are a wide array of tools, commonly used by governments to influence the economy. In this paper, we reviewed the many margins through which tax policies can affect innovation, the main driver of economic growth in the long-run. These margins include the impact of tax policy on i) the quantity and quality of innovation; ii) the geographic mobility of innovation and inventors across U.S. states and countries; iii) the declining business dynamism in the U.S., firm entry, and productivity; iv) the quality composition of firms, inventors, and teams; and v) the direction of research effort, e.g., toward applied versus basic research, or toward dirty versus clean technologies. We gave ideas drawn from research on how the right design of policy can allow policy makers to foster the most productive firms without wasting public funds on less productive ones.

The interplay between tax and innovation is arguably among the most policy-relevant and underexplored areas in endogenous growth and public finance. The scarcity of empirical studies has been due to the lack of data at the microeconomic level needed to estimate the strength of tax responses of firms and inventors. However, our computational powers are increasing at a time when many countries are making their firm- and individual-level micro data sets available to researchers. In addition, thanks to Optical Character Recognition techniques, more and more large-scale historical records are being digitized for use in economic research. These are all very exciting developments that can potentially foster this important and growing area of research.

FIGURE 1: TAXATION AND INNOVATION: FRAMEWORK

General Tax Policies

Targeted Tax Policies

