

OPTIMAL TAXATION AND R&D POLICIES

UFUK AKCIGIT

Economics, University of Chicago, NBER, and CEPR

DOUGLAS HANLEY

Economics, University of Pittsburgh

STEFANIE STANTCHEVA

Economics, Harvard, NBER, and CEPR

We study the optimal design of corporate taxation and R&D policies as a dynamic mechanism design problem with spillovers. Firms have heterogeneous research productivity, and that research productivity is private information. There are non-internalized technological spillovers across firms, but the asymmetric information prevents the government from correcting them in the first best way. We highlight that key parameters for the optimal policies are (i) the relative complementarities between observable R&D investments, unobservable R&D inputs, and firm research productivity, (ii) the dispersion and persistence of firms' research productivities, and (iii) the magnitude of technological spillovers across firms. We estimate our model using firm-level data matched to patent data and quantify the optimal policies. In the data, high research productivity firms get disproportionately higher returns to R&D investments than lower productivity firms. Very simple innovation policies, such as linear corporate taxes combined with a nonlinear R&D subsidy—which provides lower marginal subsidies at higher R&D levels—can do almost as well as the unrestricted optimal policies. Our formulas and theoretical and numerical methods are more broadly applicable to the provision of firm incentives in dynamic settings with asymmetric information and spillovers, and to firm taxation more generally.

KEYWORDS: Corporate tax, optimal taxation, subsidies, innovation, R&D.

1. INTRODUCTION

GOVERNMENTS ALL OVER THE WORLD attempt to foster innovation. In many countries, they intervene in the R&D process of private businesses through a wide variety of policies, including tax credits and deductions, direct grants and funding for research, and subsidies for R&D costs. The sheer scale of public resources spent on R&D and the variety of policies thus funded raises the question of how best to design R&D policies.

One major challenge for innovation policy is asymmetric information. 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 par-

Ufuk Akcigit: uakcigit@uchicago.edu

Douglas Hanley: doughanley@pitt.edu

Stefanie Stantcheva: sstantcheva@fas.harvard.edu

We thank Nicholas Bloom, Mike Golosov, Austan Goolsbee, Roger Gordon, Pete Klenow, Henrik Kleven, Narayana Kocherlakota, Benjamin B. Lockwood, Yena Park, Alessandro Pavan, Nicolas Serrano-Velarde, Christopher Sleet, Chad Syverson, John Van Reenen, Matthew Weinzierl, Nicolas Werquin, and numerous conference and seminar participants for feedback and comments. We thank Leo Aparisi De Lannoy, Jessica Liu, Sanjay P. Misra, and Raphael Raux for excellent research assistance. Stantcheva gratefully acknowledges the Pershing Square Foundation and the Foundations for Human Behavior Initiative for financial support. Akcigit gratefully acknowledges the National Science Foundation, the Alfred P. Sloan Foundation, and the Ewing Marion Kauffman Foundation for financial support.

ties, including the government, to observe.¹ One approach for addressing this 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 to large-scale government policies. The innovation literature has extensively addressed how to deal with spillovers from innovation, but it has not focused as much on asymmetric information about firms and how to distinguish between firms that are good at innovation and those that are not.

In this paper, we study the optimal design of taxation and R&D policies under asymmetric information. We use new methods from the public economics literature, theoretical advances in mechanism design, and firm-level data matched to patent data to discipline and quantify our analysis. We build a framework that captures this essential aspect of asymmetric information in innovation and addresses the following questions both theoretically and quantitatively: Without restricting the set of policy tools a priori, what are the best policies for promoting innovation? What key parameters do optimal policies depend on? Are there simple policies that are almost as good as the unrestricted optimal ones?

In our setting, there are two market failures that leave scope for some form of government intervention: First, there are technology spillovers between firms, whereby one firm's innovations affect other firms' productivities. Second, innovation is not appropriable and, absent Intellectual Property Rights (IPR) policy, any firm could use an "idea" embodied in an innovation. However, IPR policy may create a distortion, as is the case, for instance, of a patent system that grants firms monopoly rights.

The main impediment to fixing these market frictions in a non-distortionary way—and the key feature of our analysis—is asymmetric information. 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 (we call them "R&D investment"), others are unobservable ("R&D effort"). The firm's research productivity evolves stochastically over time. Although the firm has some information about its future productivity, it cannot perfectly foresee it. As a result, when the firm invests resources in R&D, the innovation outcomes stemming from these investments are uncertain.

In a world without private information, the government could perfectly correct for the technology externality through a Pigouvian subsidy, and for the lack of appropriability of innovation through a prize system. The asymmetric information means that the government needs to take incentive constraints into account when designing its innovation policies and limits how close the economy can get to full efficiency. We show that the need to screen firms may starkly modify the recommendations that arise with observable firm types.

Studying optimal policy under asymmetric information in a dynamic R&D investment model with spillovers is technically involved: the tractable model presented in Section 2 is one of our contributions. We pose the problem as one of mechanism design, in which we do not ex ante restrict the policies that the government can use: in this direct revelation mechanism, the government can directly choose allocations for each firm type, subject only to the asymmetric information incentive constraints. We build on new mechanism design methods described below and extend them by offering a new approach to allow

¹As shown in the empirical literature, reviewed in Section 2.4.

for spillovers between agents (firms) in the presence of asymmetric information. By doing so, we provide an entirely new and general framework to study the taxation of firms that captures key elements such as market power, investments, production, heterogeneity in productivity, intellectual property, and asymmetric information.

We first characterize the constrained efficient allocations that arise in this direct revelation mechanism with spillovers. The optimal 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 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 unobservable firm research productivity, the more rents a firm can extract if R&D investment is subsidized. This complementarity puts a brake on how well the government can correct for the technological spillovers and 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 employ more of the unobservable input, which makes the optimal R&D subsidies larger. The persistence of firms' research productivity shocks and the strength of spillovers are other key determinants of the optimal policies. We show that these constrained efficient allocations can be implemented with a parsimonious corporate income tax function.

We estimate our model using firm-level data matched to U.S. Patent Office Patent data. This allows us to measure firms' inputs into R&D, their production decisions, and their innovation output, as captured by their patents and citations. Our parameter estimates allow us to quantify the optimal policies. We can also study how well simpler innovation policies can approximate the unrestricted mechanism by comparing the revenue raised from the optimal policies to the revenue raised under restricted (and simpler) policies.

In the data, we find that R&D investments are highly complementary to a firm's research productivity and that higher productivity firms generate disproportionately more innovation from a given R&D investment. Since higher productivity firms have a comparative advantage at innovation, it is better to incentivize R&D investments less for lower productivity firms. Otherwise, it becomes excessively attractive for high productivity firms to pretend to be low productivity ones (i.e., "to mimic" low productivity firms). We discuss how these incentives translate into "wedges" and then into actual taxes and subsidies.²

Regarding the wedges, on balance, a higher net incentive for R&D for higher research productivity firms is provided with a lower profit wedge at higher profit levels and a lower R&D wedge at higher R&D levels. Intuitively, higher productivity firms are able to generate more profits from the same research investments, and an allocation with a lower profit wedge and a lower R&D wedge is more attractive to high productivity firms than to low productivity firms.

Regarding taxes and subsidies, a nonlinear, separable Heathcote–Storesletten–Violante (HSV) type subsidy combined with an HSV-type profit tax performs almost as well as the optimal policy. It features decreasing marginal profit taxes (increasing marginal profit subsidies) at higher profit levels, and decreasing marginal R&D subsidies at higher R&D investment levels. This policy perfectly mimics the shape of the wedges. Quantitatively, the most important feature is the nonlinearity in the R&D subsidy: making the profit tax

²Wedges measure the distortion in allocations relative to the laissez-faire economy's allocations and are thus akin to implicit taxes and subsidies.

linear (and lower) only generates a small welfare loss. The intuition is that a constant profit tax that is more generous than it should be for low profit firms, and at about the right level for high profit firms, does reasonably well since the loss from being too generous to low profit firms is small (because taxing their low profit levels does not yield much revenues to start with). Thus, linear corporate income taxes—common in practice—can be very close to optimal for innovating firms if combined with the right nonlinear R&D subsidy.

Related Literature. There is a longstanding static contract theory literature on the regulation of firms under private information (Laffont and Tirole (1986), Baron and Myerson (1982)). Very few papers consider the regulation of research and innovation: Sappington (1982) does so in a simple static model.

Some papers study the corrective role of personal income taxes when there are externalities such as rents (Rothschild and Scheuer (2016), Piketty, Saez, and Stantcheva (2014), Lockwood, Nathanson, and Glen Weyl (2017)). These models are static, focus on individuals rather than firms, and consider a relatively blunt tool (income taxation) because the externality-inducing action cannot be directly taxed or subsidized.

We also contribute to the new dynamic public finance literature that uses mechanism design tools to study the dynamic income taxation of agents under idiosyncratic risk. Methodologically related papers are, among others, Albanesi and Sleet (2006), Farhi and Werning (2013), Golosov, Tsyvinski, and Werning (2006), Golosov, Tsyvinski, and Werquin (2014), Sachs, Tsyvinski, and Werquin (2020), and Werquin (2016). Closest are the papers by Stantcheva (2015) and Stantcheva (2017), which incorporate endogenous investments in human capital into the (personal) dynamic tax problem. We build on the mechanism design methodology developed in Pavan, Segal, and Toikka (2014), which we augment with dynamic spillovers and a realistic, infinite-horizon dynamic life-cycle model of innovating firms, with technology spillovers. In our model, the firm's asymmetric information about its research productivity evolves stochastically over time. We also take into account the private market between intermediate and final goods producers. To solve the model with spillovers, we extend the two-step approach with an "inner" and "outer" problem proposed by Rothschild and Scheuer (2013) to this dynamic, infinite-horizon firm setting.

Theoretically, our contributions are, first, the addition of spillovers between agents (in our case, firms). Because of this important extension, the solution methods are different, both theoretically and computationally. We are able to capture key elements such as market power, investments, production, heterogeneity in productivity, intellectual property, asymmetric information, and an infinite horizon. This framework is very malleable: we illustrate several possible extensions in Appendix OA.3 of the Supplemental Material (Akcigit, Hanley, and Stantcheva (2022)) and, depending on the question at hand, parts of it can also be shut off. In particular, our model could be used to study firm taxation more broadly, when the main goal is not to incentivize innovation, but when firms have unobservable and stochastic productivity types. Computationally, we take the major step of fully estimating this dynamic model with spillovers in the data, giving precise empirical content to the variables in our model thanks to the match between patent data and firm-level data.

Grossman, Steger, and Trimborn (2013) studied the optimal time path of R&D subsidies in a standard semi-endogenous growth model and the welfare loss from implementing the long-run optimal invariant policy. There are several key differences to our setting: The authors adopted a Ramsey approach (linear policies) where they parameterized the

policies *ex ante* and had numerical solutions. We adopt a mechanism design approach. Their model contains neither heterogeneous firm productivities nor private information about these productivities.

We also use findings from the empirical literature on R&D and productivity to discipline our model and estimation (Goolsbee (1998), Bloom, Griffith, and Van Reenen (2002), Bloom and Griffith (2001), Bloom et al. (2002)). Bloom and Van Reenen (2007), Bloom, Sadun, and Van Reenen (2012), and Bloom et al. (2013) lend support to the idea that firms are heterogeneous in terms of the efficiency with which they can put their resources to productive use, and that these differences may be exceedingly difficult for the government or regulator to see. Several papers document the gap between the private and social returns to R&D and spillovers (Jones and Williams (1998, 2000)); we rely on the estimates from Bloom, Schankerman, and Van Reenen (2013) to pin down the magnitude of spillovers.

The rest of the paper is organized as follows. Section 2 presents our dynamic model and discusses its assumptions, providing empirical justification for our focus on asymmetric information in innovation. Section 3 sets up and solves the full dynamic model; Section 4 discusses the forces that shape the optimal policies. Section 5 estimates the model using firm data matched to patent data and simulates the optimal policies. Section 6 considers the welfare loss from simpler, restricted policies relative to the full optimum. Section 7 points to directions for future research. The Supplemental Material contains all proofs and a description of our computational procedure. A Supplementary Materials Appendix, within the replication file (Akcigit, Hanley, and Stantcheva (2022)), contains a simpler two-type, one-period version of the model and many sensitivity and robustness checks for the estimation.

2. A DYNAMIC MODEL OF R&D INVESTMENTS

We present a dynamic model of R&D investments with spillovers that is tractable enough to theoretically study optimal mechanism design with asymmetric information. As mentioned in the Introduction, first, this model can be used to study other types of firm investments with asymmetric information and spillovers. Second, with our core setup and methodology in place, we can incorporate additional aspects of R&D investments by firms. Some of these generalizations are discussed in Section 2.5, together with our modeling choices. Finally, by turning off certain aspects such as spillovers and specifying a particular market structure between final and intermediate goods producers, our framework is also amenable to studying firm taxation with heterogeneous firms more generally, even for non-innovating firms.

2.1. Setting

At the core of the model are firms, producing and selling differentiated intermediate goods. They engage in R&D to improve the quality of their differentiated products through innovation. There are both observable and unobservable R&D inputs. More precisely, the quality q_t at time t of the intermediate good evolves according to

$$q_t = H(q_{t-1}, \lambda_t),$$

where λ_t is the endogenous quality improvement for period t , which we call the “step size”:

$$\lambda_t = \lambda_t(r_{t-1}, l_t, \theta_t).$$

The step size depends on three components:

- (i) *Observable R&D inputs*: r_{t-1} denotes the resources that the firm spent on R&D in period $t - 1$. They include the pay of scientists and researchers, lab equipment, material supplies, and raw materials for research and innovation. Their monetary cost is $M_t(r_t)$, with $M'_t(r_t) > 0$ and $M''_t(r_t) \geq 0$.³ We call these observable inputs “R&D investments.”
- (ii) *Unobservable R&D inputs*: Each firm also needs to provide some unobservable R&D inputs, which cannot be directly monitored by the government. One such input is unobservable research effort, which is required in order to transform the material resources into an innovation output. We call these unobservable R&D inputs “R&D effort” for concreteness, although they could include other costly, unobservable actions taken at the organizational level that contribute to research. These unobservable R&D inputs are denoted by l_t and entail a cost $\phi_t(l_t)$ for the firm, which is increasing and convex.
- (iii) *Firm type*: Every firm has a type θ_t that determines the efficiency with which it converts the observable and unobservable inputs r_{t-1} and l_t into innovation (product quality), called “research productivity.” For instance, θ may represent the efficiency of management, an interpretation bolstered by recent papers on the importance and heterogeneity of management practices across firms (Bloom and Van Reenen (2007), Bloom, Sadun, and Van Reenen (2012), Bloom et al. (2013)). The type can also be a composite measure of several exogenous characteristics of a firm that shape its efficiency in producing innovations, such as the quality of its organization, of its business model, or of its “ideas.” What is key is that firms differ in their ability to produce innovation and that this ability is hard to observe by a government or regulator.

It is critical to bear in mind that, for policy design purposes, it is equivalent whether a characteristic (such as research productivity) is truly unobservable or whether it is simply impossible to condition policies on it. In either case, it is necessary to include the incentive compatibility constraints that will be at the core of our mechanism design problem. In addition, if innovation output depends on unobservable (or non-verifiable) characteristics such as research productivity, it also means that there are some unobservable inputs. These unobservable inputs prevent the government from perfectly “inverting” the innovation outcome (conditional on observable inputs) to obtain the firm’s productivity type. If all inputs were perfectly observable (i.e., the firm could not misreport them), there would be no asymmetric information problem. In Section 2.4, we provide abundant empirical evidence on the prevalence of asymmetric information in the innovation arena.

The type θ_t evolves over time according to a Markov process $f^t(\theta_t|\theta_{t-1})$ on $\Theta = [\underline{\theta}, \bar{\theta}]$. Denote by θ^t the history of type realizations until time t , that is, $\theta^t = \{\theta_1, \dots, \theta_t\}$, and by $P(\theta^t)$ the probability of that history:

$$P(\theta^t) := f^t(\theta_t|\theta_{t-1}) \dots f^1(\theta_1).$$

We assume that

$$\frac{\partial \lambda}{\partial \theta} > 0, \quad \frac{\partial \lambda}{\partial r} > 0, \quad \frac{\partial \lambda}{\partial l} > 0, \quad \text{and} \quad \frac{\partial^2 \lambda}{\partial \theta \partial l} > 0,$$

³Taking a broad view of these material inputs is consistent with the fact that many types of material inputs and expenses are eligible for R&D tax credits or subsidies (Tyson and Linden (2012)).

so that higher realizations of research productivity θ , higher R&D investments, and higher effort lead to a larger step size, and the marginal returns to effort are higher for higher types of firms (this assumption will permit screening types).

Let us emphasize the two related, but conceptually very distinct terms used: Firms' product *quality* refers to the product quality q_t of the intermediate good produced by the firm. Firm *research productivity* refers to the efficiency type of the firm, θ_t , which affects the innovation process that produces the product quality q_t .⁴

Note that because the step size depends on lagged R&D investments and on the stochastic realization of θ_t , about which the firm has some, but not perfect, advance information at the time the R&D investment decisions r_{t-1} are made, the returns to R&D are both stochastic and heterogeneous across different types of firms. This captures the notion that spending on R&D has uncertain returns and is not guaranteed to lead to a good innovation. The distinction in the timing between R&D investments and effort has no technical implications and will not change our results qualitatively or quantitatively.⁵ Conceptually, R&D investments can be thought of as observable investments that—much like physical capital investments—take a while to yield returns and are determined before the uncertainty is realized. R&D effort can be viewed as inputs that can more easily be adjusted in response to the current state, that is, utilization rate of the equipment, managerial input, process improvements, labor effort of researchers, etc.

Input Complementarity. We can characterize the complementarity between the three different inputs that enter the step size using the Hicksian coefficient of complementarity (Hicks (1970)), which will be important for our results. For any two variables $(x, y) \in \{\theta_t, r_{t-1}, l_t\} \times \{\theta_t, r_{t-1}, l_t\}$, the Hicksian coefficient of complementarity between variables x and y in the step size creation is denoted by

$$\rho_{xy} = \frac{\frac{\partial^2 \lambda}{\partial x \partial y} \lambda}{\frac{\partial \lambda}{\partial x} \frac{\partial \lambda}{\partial y}}.$$

The higher coefficient ρ_{xy} is, the more inputs x and y are complementary in the production of the step size. To give a few examples, suppose that the step size function takes the multiplicatively separable form

$$\lambda_t(r_{t-1}, l_t, \theta_t) = h_t^1(r_{t-1})h_t^2(l_t)h_t^3(\theta_t)$$

for some increasing functions h_t^1 , h_t^2 , and h_t^3 . Then, $\rho_{\theta l} = \rho_{\theta r} = \rho_{lr} = 1$. On the other hand, an additively separable step size function

$$\lambda_t(r_{t-1}, l_t, \theta_t) = h_t^1(r_{t-1}) + h_t^2(l_t) + h_t^3(\theta_t)$$

would have $\rho_{\theta l} = \rho_{\theta r} = \rho_{lr} = 0$. Finally, a CES function of the form

$$\lambda_t(r_{t-1}, l_t, \theta_t) = (\alpha_r r_{t-1}^{1-\rho_t} + \alpha_\theta \theta_t^{1-\rho_t} + \alpha_l l_t^{1-\rho_t})^{\frac{1}{1-\rho_t}}$$

has $\rho_{\theta l} = \rho_{\theta r} = \rho_{lr} = \rho_t$.

⁴To clarify a sometimes confusing point: once produced, innovations are non-rival and non-appropriable absent IPR. The inputs into that innovation are, as usual, rival.

⁵If the timing was contemporaneous, the sums and expectations for the R&D wedge in Proposition 1 should simply be lagged one period.

Quality Spillovers. An important element of the model is the presence of spillovers between firms. One firm’s innovation has a beneficial effect on the production costs of other firms. Such spillovers can reflect the direct use of better technologies and processes in production and learning from new technologies to improve one’s production. The specific shape of the knowledge spillovers in our model is taken from Akcigit and Kerr (2018) to capture the idea of “building on the shoulders of giants” (Aghion and Howitt (1992), Romer (1990)). Importantly, however, the exact shape of the spillovers is not key for our theoretical results and the spillovers could appear in different parts of the model, as discussed in more detail below.⁶ Aggregate quality is given by

$$\bar{q}_t = \int_{\Theta^t} q_t(\theta^t) P(\theta^t) d\theta^t.$$

The production cost of each firm is decreasing in aggregate quality so that the cost of producing k units of intermediate goods is $C_t(k, \bar{q}_t)$.

Final Goods Production. The final good is consumed by consumers and is produced competitively using the intermediate goods as inputs. The production technology for the final good is

$$Y_t = \int_{\Theta^t} Y(q_t(\theta^t), k_t(\theta^t)) P(\theta^t) d(\theta^t),$$

where $Y(q_t(\theta^t), k_t(\theta^t))$ is the contribution of the intermediate good of firm θ^t to the final good, and depends on the quantity $k_t(\theta^t)$ and the quality $q_t(\theta^t)$ of the intermediate good of firm θ^t . The price of the final good is normalized to 1. The demand function for the intermediate good that arises in the market will depend on the IPR regime.

Patent Protection and Monopoly Power. In this setting, one way of capturing different IPR regimes is through different demand functions $p(q_t(\theta^t), k_t(\theta^t))$. Our benchmark case mirrors the current state of the world and grants the innovating firm full patent protection. Thus, the intermediate goods producer has monopoly power and faces a downward sloping demand curve derived from the optimization problem of the final goods producer, which is a function of the quality and quantity, $p(q, k) = \frac{\partial Y(q,k)}{\partial k}$.

Firm Life Cycle. Firms live for an infinite number of periods. We assume a small open economy with gross interest rate R . Let $\theta^t|\theta_1$ denote a history θ^t such that the period 1 type realization is θ_1 and let $P(\theta^t|\theta_1)$ be the probability of that history after initial realization θ_1 . In the laissez-faire economy, the firm chooses quality $q_t(\theta^t)$, quantity $k_t(\theta^t)$, R&D investments $r_t(\theta^t)$, and R&D effort $l_t(\theta^t)$ to maximize its objective given its initial type θ_1 , initial quality q_0 , and R&D investments r_0 :

$$\sum_{t=1}^{\infty} \left(\frac{1}{R}\right)^{t-1} \int_{\Theta^t} (p(q_t(\theta^t), k_t(\theta^t))k_t(\theta^t) - C(k_t(\theta^t), \bar{q}_t) - M_t(r_t(\theta^t)) - \phi_t(l_t(\theta^t))) \times P(\theta^t|\theta_1) d(\theta^t|\theta_1) \tag{1}$$

⁶In brief, all our formulas will be expressed at a general level as functions of net output and profits, which will depend on own quality and aggregate quality. The channel could be through the cost (as here), directly through the demand function (note that the equilibrium price always depends on aggregate quality), or through the innovation production function.

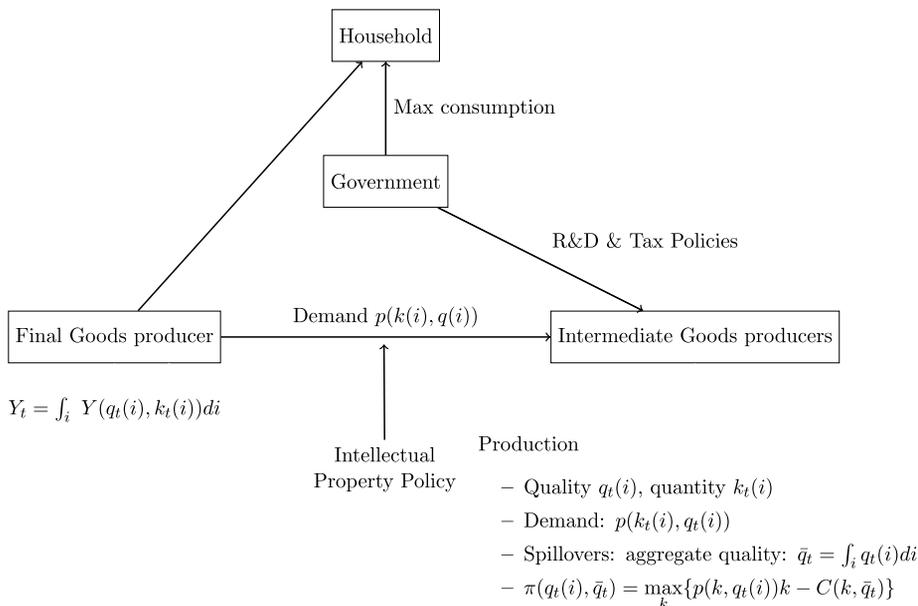


FIGURE 1.—Model summary.

subject to the law of motion of quality $q_t(\theta^t) = H(q_{t-1}(\theta^{t-1}), \lambda_t(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t))$.

Production Decision. Given the demand function $p(q, k)$, let production profits gross of R&D costs be

$$\pi(q_t(\theta^t), \bar{q}_t) := \max_k \{p(q_t(\theta^t), k)k - C(k, \bar{q}_t)\}.$$

The firm’s maximization pins down the quantity produced for a given quality level. Figure 1 summarizes the model in schematic and static form.

2.2. Social Welfare

Consumer surplus is equal to the consumption of the final good, net of all transfers to firms. The gross transfer to the firm of type θ^t in period t is the sum of its production costs ($C(k_t(\theta^t), \bar{q}_t)$), R&D costs ($M_t(r_t(\theta^t))$), and a net transfer denoted by $T_t(\theta^t)$. The exact shape of this net transfer will be specified depending on the market structure and information structure in each of the cases considered below (in the laissez-faire case, the gross transfer is just price times quantity and the firm payoff is as in (1)). Consumer surplus in period t is thus: $Y(k_t(\theta^t), q_t(\theta^t)) - (C(k_t(\theta^t), \bar{q}_t) + M_t(r_t(\theta^t)) + T_t(\theta^t))$. Let $v_t(\theta^t)$ be the period t payoff (surplus) of a firm with history θ^t :

$$v_t(\theta^t) = T_t(\theta^t) - \phi_t(l_t(\theta^t)). \tag{2}$$

Social welfare (the objective the planner maximizes) is a weighted sum of consumer surplus plus firm surplus:⁷

$$\sum_{t=1}^{\infty} \left(\frac{1}{R}\right)^{t-1} \left(\int_{\theta^t} (Y(k_t(\theta^t), q_t(\theta^t)) - C(k_t(\theta^t), \bar{q}_t) + M_t(r_t(\theta^t)) + T_t(\theta^t)) + (1 - \chi)v_t(\theta^t))P(\theta^t) d(\theta^t) \right). \tag{3}$$

The key benchmark case in the contract theory literature has $\chi = 1$ so that the social objective becomes maximizing total social surplus (consumer plus firm surplus), minus all informational rents, the so-called “virtual surplus.” Note also that, even absent any redistributive concerns, maximizing efficiency essentially amounts to maximizing a weighted sum of surpluses of consumers and firms, if we assume, as is standard in the contract theory literature, that the planner can only raise the money for transfers through some distortionary method (e.g., excise taxes or distortionary income taxes on households), so that the cost of one unit of transfer is weakly greater than 1 (see [Laffont and Tirole \(1986\)](#)).

2.3. Two Market Failures and First-Best Allocation

There are two market failures in this setting (in the absence of any government intervention): First, the lack of appropriability of innovation means that there will be no investment in innovation as long as producers’ profits are not protected by some IPR. Second, there are non-internalized technology spillovers that affect others’ production technologies.

Suppose the planner could observe firm types and that transfers are perfectly non-distortionary ($\chi = 0$).⁸ Social welfare is then $W^{\text{first-best}}$, equal to total expected discounted output net of production costs, R&D investment costs, and R&D effort costs:

$$W^{\text{first-best}} = \sum_{t=1}^{\infty} \left(\frac{1}{R}\right)^{t-1} \left(\int_{\theta^t} (Y(k_t(\theta^t), q_t(\theta^t)) - C(k_t(\theta^t), \bar{q}_t) - M_t(r_t(\theta^t)) - \phi_t(l_t(\theta^t)))P(\theta^t) d(\theta^t) \right).$$

The first-best maximization program is

$$\max_{\{l_t(\theta^t), r_t(\theta^t), k_t(\theta^t)\}_{t, \theta^t}} W^{\text{first-best}} \quad \text{s.t. } q_t(\theta^t) = H(q_{t-1}(\theta^{t-1}), \lambda_t(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t))$$

with q_0 and r_0 given.

Conditional on a given quality $q_t(\theta^t)$, the production choice of the planner is $k^*(q_t(\theta^t), \bar{q}_t)$. Denote by $Y^*(q_t(\theta^t), \bar{q}_t) = Y(k_t^*(q_t(\theta^t), \bar{q}_t), q_t(\theta^t))$ the optimized consumption of the intermediate good, and by $\tilde{Y}^*(q_t(\theta^t), \bar{q}_t) = Y^*(q_t(\theta^t), \bar{q}_t) - C(k_t^*(q_t(\theta^t), \bar{q}_t), \bar{q}_t)$ consumption net of production costs for the intermediate good.

For the exposition, we simplify the accumulation equation of quality to be

$$q_t = (1 - \delta)q_{t-1} + \lambda_t \quad \text{with } 0 < \delta < 1,$$

⁷The final goods producer always has zero payoff because it operates under perfect competition.

⁸Under full information, type-specific lump-sum transfers and taxes are feasible.

where δ is the depreciation factor. None of the results depend on this simplification, but the notation is much lighter.

Firms then choose R&D investment and effort so that their total marginal *social* benefit equals their marginal costs:

$$M'_t(r_t(\theta^t)) = \frac{1}{R} \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \left(\frac{\partial \tilde{Y}^*(q_s(\theta^s), \bar{q}_s)}{\partial q_s} + \frac{\partial \tilde{Y}^*(q_s(\theta^s), \bar{q}_s)}{\partial \bar{q}_s} \right) \frac{\partial \lambda_{t+1}(\theta^{t+1})}{\partial r_t(\theta^t)} \right),$$

$$\phi'_t(l_t(\theta^t)) = \mathbb{E} \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \left(\frac{\partial \tilde{Y}^*(q_s(\theta^s), \bar{q}_s)}{\partial q_s} + \frac{\partial \tilde{Y}^*(q_s(\theta^s), \bar{q}_s)}{\partial \bar{q}_s} \right) \right) \frac{\partial \lambda_t(\theta^t)}{\partial l_t(\theta^t)},$$

where the expectation operator is over histories θ^t .

2.4. Asymmetric Information and Government Policies

Asymmetric Information Structure. The core asymmetry of information, which holds throughout this paper, is that the history of research productivity realizations θ^t and the unobservable R&D effort l_t are private information for each firm. In the benchmark case, the government observes the full histories of R&D investment r_t , quality improvements (the step size λ_t), and the realized quality q_t . To make this more concrete, think of the government observing past patents granted to each firm and their citations. Quantity $k(\theta^t)$ is unobservable as well, or, equivalently, cannot be conditioned on by the government. This amounts to saying that the government cannot intervene directly in the market between the intermediate and final goods producer and has to take as given their production decisions. In Section S.2 of the replication file, we consider the case in which the government can intervene in that market because quantity is observable.

Government Policies Considered. We study several types of government policies. First, we take a mechanism design approach and consider the optimal unrestricted direct revelation mechanism, which is subject only to the incentive compatibility constraints that arise due to asymmetric information on firm type, R&D effort, and quantity produced. We relax the unobservability of quantity in Section S.2 of the replication file. We do not constrain policy tools ex ante, but rather find the optimal allocations subject to only incentive compatibility constraints and then show what tax functions can implement these allocations (Section 4.2). We subsequently study the shape of and revenue losses from restricted, parametric instruments, which are simpler (Section 6).

The Importance of Asymmetric Information. We now highlight why asymmetric information is a crucial feature in the innovation process. First, we summarize the abundant literature showing the prevalence of asymmetric information; second, we show in our data that it is very difficult to predict a firm's innovation quality based on observables.

In our model, the productivity type of the firm, θ , embodies elements such as the quality of the manager, of its organization, business model, or ideas. It is quite clear that these elements are very hard to observe or, equivalently from the point of view of the government, to condition policies on.⁹ A large literature argues that asymmetric information

⁹Recall that for policy design purposes, it is equivalent whether a variable is truly unobservable or simply impossible to condition policies on. In both cases, the incentive compatibility constraints at the core of our mechanism design problem are needed.

is likely to be a key issue in innovation. Hall and Lerner (2010) summarized several of these contributions. In their terminology, asymmetric information refers to the fact that the innovator has “better information about the likelihood of success” than anyone else, including investors and the government. Based on the abundant literature on the asymmetric information between innovators and investors—which leads to financing frictions and inefficiencies—they argued that such informational frictions are likely to carry over in an even more pronounced way to the interaction between inventors and the government. They also cautioned against trying to reduce information asymmetry by mandating fuller disclosure, which can be entirely unproductive in the innovation arena because innovations can be easily imitated. Thus, revealing one’s productivity (quality of the idea, management style, or organizational process) to the government runs the risk of revealing it to one’s competitors, which will distort the quality of the signal provided.

The need for screening is embodied in the existence and size of the venture capitalist (VC) industry. Gompers (1995) and others have argued that VCs tend to operate in areas where asymmetric information problems are more common, such as high-technology and innovating sectors. Kaplan and Stromberg (2001) also documented the intensive efforts that VCs put into screening possible entrepreneurs in order to directly circumvent asymmetric information issues. The severity of the asymmetric information problem is illustrated by the fact that, “even highly-skilled VCs cannot distinguish in advance the next Google from the other cases” (Kerr, Nanda, and Rhodes-Kropf (2014)). Given that VCs face asymmetric information problems despite the huge time investment and detailed involvement in the firms that they fund, it is hard to imagine that the government would not be facing much larger informational problems when designing a decentralized tax system that does not micro-manage or directly intervene in firms.

Several papers have also looked at responses of stock prices as another symptom of the asymmetric information problem inherent in innovation (Zantout (1997), Alam and Walton (1995), Gharbi, Sahut, and Teulon (2014)). Aboody and Lev (2000) showed that insider gains are larger at R&D intensive firms than firms without R&D because that is where asymmetric information is higher.

Finally, the literature also highlights that not all R&D inputs are easily verifiable. Hall and Van Reenen (2000) called this the “relabeling” problem and offered many examples. Mansfield (1986) surveyed the effects of R&D tax credits in the United States, Canada, and Sweden and found that there is substantial misreporting. More recently, Chen et al. (2021) documented that around 30% of reported R&D investment by Chinese firms could be due to relabeling.

We can also directly provide some suggestive evidence for asymmetric information in our data. We study what share of the innovation quality of a firm can actually be predicted based on observables. We explain our data and measurement in more detail in Section 5. In brief, we measure the quality of the innovations of a firm by its patent citations, namely, all forward citations that accrue to a firm’s patents until today (Hall, Jaffe, and Trajtenberg (2001)). We regress the citations-weighted patents of a firm on a whole range of controls, such as sector and year fixed effects (or even the interaction between these two), lagged sales, employment, R&D spending, age, balance sheet variables, etc. We then look at how well we can predict the quality of the firm’s innovations. The prediction is quite poor. Even adding such an exhaustive list of control variables, the *R*-squared of these regressions barely moves above 0.3. In addition, it is especially difficult to predict performance based on data from the first few years of a company’s life cycle (when there is only a short track record available) and very difficult to predict which firms will become “superstars,” that is, receive highly-cited and influential patents. Again, this set of information is likely a

very generous upper bound on what the government could realistically condition taxes on. Furthermore, if taxes actually depended on these variables, firms would of course respond along these margins, too (like they do along the profit and R&D margins in this paper), so they are not tag-like signals that are immutable to taxation.¹⁰

2.5. Discussion of the Assumptions and Possible Generalizations

Additional Firm Heterogeneity. Firms may be heterogeneous along many dimensions, such as their sector or the type of product. If the government or regulator wants to fine tune the policy for firms according to some observable vector of characteristics X , then the mechanism needs to condition on X . Since X is observable, this does not require adding any incentive constraints and only increases the state space to be kept track of. We explicitly discuss heterogeneity in firms' production productivities in Section OA.3 of the Supplemental Material.

Entry and Exit. In principle, firms in our model only make intensive margin decisions about how much to produce. Exit is captured only through the discount rate R that combines the interest rate and an exogenously given exit or death rate. The discount rate could also depend on age. Regarding entry, firms in the model enter jointly with their cohort. Free entry could affect the size of a cohort and entry barriers could be studied as a policy tool in the model as well.¹¹

The Role of IPR. Our focus is not on IPR, but on the design of R&D policies. However, the shape and magnitude of optimal R&D policies depend on the IPR policies. Our starting point is to model the IPR policy as it currently is in the world, namely, granting patent protection and monopoly rights to innovating firms.¹² As a result, part of the role of R&D policies will be to partially correct for the monopoly distortion induced by the patent system.¹³ We also consider two different cases based on whether the government can intervene in the private market between intermediate and final goods producers, that is, whether it can observe and make the optimal policy contingent on the quantity k produced. Our benchmark case is when the government cannot control quantity. We cover a setting in which the government can control quantity in Section S.2 of the replication file. In this case, given that quality is observable, the government can incentivize the socially optimal quantity to be produced and thus counteract the monopoly distortion.

Shape of the Spillovers. The exact shape of the spillovers will not be important for our theoretical results and will not affect the forces we describe and the key qualitative

¹⁰These results are by no means a formal "test" of asymmetric information. It may be that the prediction could be improved with different, better data or methods and it is always very difficult to disentangle heterogeneity (generating asymmetric information) from uncertainty (which is consistent with symmetric information).

¹¹For instance, the government could endogenously set a lower bound for θ .

¹²One may instead consider another system, such as patent for protection for x years or patent protection for a fraction of the monopoly profits.

¹³If the world were different, and there was an IPR policy that did not grant monopoly power, for example, a prize system, then the R&D policies would not be set to make up for the monopoly distortion. Whenever the product quality is observable, the optimal IPR is very simple and amounts to paying the innovating intermediate goods producer a prize to buy the innovation, and then produce the socially optimal quantity. An equivalent system is to have full patent protection, but pay a nonlinear price subsidy to the monopolist that aligns the private valuation of quantity with its social valuation.

mechanisms. Following Akcigit and Kerr (2018), we suppose spillovers affect the costs of production. This captures the idea of “building on the shoulders of giants” in innovation models. Innovations improve the productivity of production labor and/or the process with which firms produce. Think, for instance, of computers (an innovation from the point of view of one or several firms) that are then being bought and used in other companies to produce better, cheaper, and faster. Alternatively, one can think of other innovations in communication technologies, production technologies, or health improvement, etc. However, our theoretical framework is general enough that spillovers could appear in other parts. We could instead specify them as directly affecting the cost of producing innovations: $q_t = H(q_{t-1}, \lambda_t, \bar{q}_t)$. The formulas below are expressed in terms of general profit functions or net output functions that depend in a reduced-form way on both own quality and aggregate quality. This is a quite general formulation that will apply for many types of spillovers, regardless of the specific functional form assumptions. Another possible variant would be to let lagged aggregate quality \bar{q}_t enter either the production cost function or the innovation production function. This will merely cause a shift in the time indices in the formulas, but not change anything substantial.

Different Types of Investments With Different Externalities. It is possible to consider different types of firm investments that each generate different externalities (see Section OA.3 of the Supplemental Material). For instance, investment in new drug discovery may have larger positive spillovers than investment aimed at improving machinery that is only used by few firms.

3. A DYNAMIC DIRECT REVELATION MECHANISM WITH SPILLOVERS

Recall that each firm’s history θ^t and research effort l_t are private information. The government observes the step size λ_t , the realized quality q_t , and the R&D investment r_t . To solve for the constrained efficient allocations, we imagine that the government designs a direct revelation mechanism in which, every period, each firm reports a type $\theta'_t(\theta^t)$ as a function of their history θ^t . Denote a reporting strategy by $\sigma = \{\theta'_t(\theta^t)\}_{t=1}^\infty$. A reporting strategy generates a history of reports $\theta'^t(\theta^t)$. The government then assigns allocations of step sizes and R&D investments, denoted by $x(\theta'^t) = \{\lambda(\theta'^t), r(\theta'^t)\}_{\theta^t}$ and a transfer $T_t(\theta'^t)$ as functions of the history of reports. For simplicity, we normalize the starting R&D investment for all agents to be $r(\theta^0) = r_0$.¹⁴ Let $l_t(\lambda_t(\theta'^t(\theta^t)), r(\theta'^{t-1}(\theta^{t-1}), \theta_t))$ denote the R&D effort that would have to be provided for true type θ_t who reports θ'^t (and, hence, had to invest $r(\theta'^{t-1}(\theta^{t-1}))$ in the previous period and has to produce a step size of $\lambda_t(\theta'^t(\theta^t))$). We can make the following assumption for simplicity.

ASSUMPTION 1: (l_t, r_t) belongs to a convex and compact set.

Suppose that the vector of aggregate qualities $\{\bar{q}_t\}_{t=1}^\infty$ is given. The continuation value after history θ^t under reporting strategy σ , denoted by $V^\sigma(\theta^t)$, given allocation rule is

$$V^\sigma(\theta^t) = T_t(\theta'^t(\theta^t)) - \phi_t(l_t(\lambda_t(\theta'^t(\theta^t)), r_t(\theta'^{t-1}(\theta^{t-1}), \theta_t))) + \frac{1}{R} \int_{\Theta} V^\sigma(\theta^{t+1}) f^{t+1}(\theta_{t+1}|\theta_t) d\theta_{t+1}.$$

¹⁴Since r_0 is observable, if it were heterogeneous across firms, allocations would need to be specified as functions of (θ^t, r_0) , which does not complicate the problem, but makes the notation heavier.

$V^\sigma(\theta^t)$ depends on the report-contingent allocations specified by the government, but this dependence is implicit to lighten the notation. Let the continuation value under truthful reporting be $V(\theta^t)$. Incentive compatibility requires that, after every history, and for all reporting strategies σ ,

$$V(\theta^t) \geq V^\sigma(\theta^t) \quad \forall \sigma, \theta^t.$$

Under truth-telling, the continuation utility as of the first period in sequential form is

$$V_1(\{\lambda_s(\theta^s), r_s(\theta^s), T_s(\theta^s)\}_{s=1}^\infty, \theta_1) = \sum_{t=1}^\infty \left(\frac{1}{R}\right)^{t-1} \cdot \left\{ \int_{\theta^t} \{T_t(\theta^t) - \phi_t(l_t(\theta^t))\} P(\theta^t | \theta_1) d\theta^t \right\}$$

with $l_t(\theta^t) := l_t(\lambda_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t)$. (4)

3.1. A First-Order Approach

We use a first-order approach, which replaces all the incentive constraints of agents with their envelope conditions.¹⁵ If the agent's report after history θ^t is optimally chosen, the envelope theorem tells us that the change in continuation utility from a change in the type is only equal to the direct effect of the type on utility (the indirect effect of the type on the allocation through the report is zero).

We now focus on a Markov process, although many of the results are generalizable to a broader set of processes. Let $I_{1,t}(\theta^t)$ be the impulse response function of the type realization in period t to a shock in the type realization at time 1, defined as, for any Markov process,

$$I_{1,t}(\theta^t) = \prod_{s=2}^t \left(- \frac{\frac{\partial F^s(\theta_s | \theta_{s-1})}{\partial \theta_{s-1}}}{f^s(\theta_s | \theta_{s-1})} \right).$$

The impulse response function captures the persistence of the stochastic type process. For instance, for an autoregressive process where $\theta_t = \tilde{\rho}\theta_{t-1} + \varepsilon_t$, the impulse response is simply $I_{1,t}(\theta^t) = \tilde{\rho}^{t-1}$. We now make two technical assumptions that will allow us to apply the first-order approach, and which are directly adapted from [Milgrom and Segal \(2002\)](#).

ASSUMPTION 2: $f^s(\theta_s | \theta_{s-1}) > 0 \quad \forall \theta_s, \theta_{s-1} \in \Theta$.

This is the full support assumption, which can be relaxed as in [Farhi and Werning \(2013\)](#) to allow for a moving support over time.

ASSUMPTION 3: $\frac{\partial F^s(\theta_s | \theta_{s-1})}{\partial \theta_{s-1}}$ exists, is bounded, and $\frac{\partial F^s(\theta_s | \theta_{s-1})}{\partial \theta_{s-1}} \leq 0$.

Assumption 3 states that the distribution function is differentiable in θ_{t-1} , that its derivative is bounded, and that a higher type realization in period s increases the realization of the period $s+1$ type in a first-order stochastic dominance sense. If it is satisfied, then $I_{s,t}(\theta^t)$ is well-defined, non-negative, and bounded. We could replace the boundedness assumptions with the assumption that $F^s(\theta_s | \theta_{s-1})$ is either convex or concave in θ_{s-1} on Θ . All the examples we discuss, such as an AR(1), log AR(1), i.i.d., or a fully persistent process, satisfy this assumption.

¹⁵See [Pavan, Segal, and Toikka \(2014\)](#), [Farhi and Werning \(2013\)](#), and [Stantcheva \(2017\)](#).

We can rewrite the per-period payoff of the firm from (2) as a function of the allocation of transfer, step size, and past R&D spending and given its true type θ_t :

$$v_t(T_t, \lambda_t, r_{t-1}; \theta_t) = T_t - \phi_t(l_t(\lambda_t, r_{t-1}, \theta_t)).$$

Note that $\frac{\partial v_s}{\partial \theta_s} = \phi'_s(l_s(\theta^s)) \frac{\partial \lambda_s(l_s(\theta^s), r_{s-1}(\theta^{s-1}), \theta_s) / \partial \theta_s}{\partial \lambda_s(l_s(\theta^s), r_{s-1}(\theta^{s-1}), \theta_s) / \partial \lambda_s}$. Because of Assumption 1 and the continuity of ϕ'_s and λ_s , this expression is bounded. The envelope condition in its derivative form is given by

$$\frac{\partial V(\theta^t)}{\partial \theta_t} = \mathbb{E} \left(\sum_{s=t}^{\infty} I_{t,s}(\theta^s) \left(\frac{1}{R} \right)^{s-t} \frac{\partial v_s(\theta^s)}{\partial \theta_s} \mid \theta_t \right). \tag{5}$$

Let $V_1(\theta_1)$ be the expected continuation utility as of period 1 for agents with initial type θ_1 . The participation constraints are, for all θ_1 ,

$$V_1(\theta_1) \geq 0. \tag{6}$$

The integral form of this envelope condition at history θ^t is

$$V(\theta^{t-1}, \theta_t) = \int_{\underline{\theta}}^{\theta_t} \frac{\partial V(\theta^{t-1}, m)}{\partial m} dm + V(\theta^{t-1}, \underline{\theta}). \tag{7}$$

This gives an expression for the informational rent the principal must give to the agent at node θ^t to entice the agent to report their true type.

3.2. Planner’s Problem

The planner’s objective is to maximize social welfare in (3) subject to the incentive constraints in (5) and participation constraints in (6). For simplicity, we set $\chi = 1$.¹⁶

Fix a given sequence of aggregate qualities, $\bar{q} = \{\bar{q}_1, \dots, \bar{q}_T\}$. The planner cannot directly choose the quantity, so the intermediate goods producer will choose its quantity $k(q_t(\theta^t), \bar{q}_t)$ to maximize profits $p(q_t(\theta^t), k)k - C(k, \bar{q}_t)$. This yields consumption net of production costs equal to $\tilde{Y}(q_t(\theta^t), \bar{q}_t) = Y(q_t(\theta^t), k(q_t(\theta^t), \bar{q}_t)) - C(k(q_t(\theta^t), \bar{q}_t), \bar{q}_t)$. The objective becomes

$$W(\bar{q}) = \mathbb{E} \left\{ \sum_{t=1}^{\infty} \left(\frac{1}{R} \right)^{t-1} \{ \tilde{Y}(q_t(\theta^t), \bar{q}_t) - M_t(r(\theta^t)) - T_t(\theta^t) \} \right\}.$$

Using the expression for $V_1(\theta_1)$ from (4), we can replace the sum of transfers $T_t(\theta^t)$ to obtain

$$-\mathbb{E} \left(\sum_{t=1}^{\infty} \left(\frac{1}{R} \right)^{t-1} T_t(\theta^t) \mid \theta_1 \right) = -V_1(\theta_1) - \mathbb{E} \left(\sum_{t=1}^{\infty} \left(\frac{1}{R} \right)^{t-1} \phi_t(l_t(\theta^t)) \mid \theta_1 \right).$$

Under Assumption 3, all that is needed to satisfy all participation constraints is to set $V_1(\underline{\theta}_1) = 0$. Using the expression for the informational rent that needs to be forfeited to

¹⁶This is the typical case in the contract theory literature, which aims to maximize total social surplus (efficiency) and minimize rents. Any $\chi < 1$ will simply appear as a scaling factor in front of the “screening term” in all formulas below.

each agent from (7), the expected discounted payoff to the planner is the “virtual surplus,” that is, the total social surplus minus informational rents:

$$W(\bar{q}) = \mathbb{E} \left\{ \sum_{t=1}^{\infty} \left(\frac{1}{R} \right)^{t-1} \left\{ \tilde{Y}(q_t(\theta^t), \bar{q}_t) - M_t(r(\theta^t)) - \phi_t(l_t(\theta^t)) - \frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t} \frac{\partial v_t(\theta^t)}{\partial \theta_t} \right\} \right\}.$$

The planner’s problem can be split into two steps. In the first step, called the “partial” problem, the sequence of aggregate qualities $\bar{q} = \{\bar{q}_1, \dots, \bar{q}_T\}$ is taken as given. The planner solves for the optimal allocations subject to resource and incentive constraints as functions of this conjectured sequence. To ensure that the sum of aggregate qualities that arises is consistent with the conjectured \bar{q} , the planner must account for a consistency constraint in each period t :

$$\int_{\Theta^t} q_t(\theta^t) P(\theta^t) d\theta^t = \bar{q}_t.$$

Let $\frac{\eta_t}{R^{t-1}}$ be the multiplier on the consistency constraint in period t . The maximum of this problem is denoted by $P(\bar{q})$.

Partial problem: The program for a given sequence \bar{q} is to choose $\{\lambda_t(\theta^t), l_t(\theta^t), r_t(\theta^t)\}_{\Theta^t}$ so as to solve

$$P(\bar{q}) = \max W(\bar{q}) \quad \text{s.t.}$$

$$\int_{\Theta^t} q_t(\theta^t) P(\theta^t) d\theta^t = \bar{q}_t \quad \text{and} \quad q_t(\theta^t) = q_{t-1}(\theta^{t-1})(1 - \delta) + \lambda_t(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t).$$

Using the expression for $\frac{\partial v_t}{\partial \theta_t}$, we have

$$W(\bar{q}) = \sum_{t=1}^{\infty} \left(\frac{1}{R} \right)^{t-1} \left\{ \int_{\Theta^t} \left\{ \tilde{Y}(q_t(\theta^t), \bar{q}_t) - M_t(r(\theta^t)) - \phi_t(l_t(\theta^t)) - \frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t} \left[\phi'_t(l_t(\theta^t)) \frac{\partial \lambda_t(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t)}{\partial \lambda_t(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t)} \right] \right\} P(\theta^t) d\theta^t \right\}.$$

Full problem: The full program consists in optimally choosing the sequence \bar{q} , given the values $P(\bar{q})$ solved for in the first step:

$$P : \max_{\bar{q}} P(\bar{q}). \tag{8}$$

Verifying Global Incentive Constraints. Since the first-order approach is built on only necessary (but not necessary and sufficient) conditions, we need to perform a numerical ex post verification to check that the allocations found are indeed (globally) incentive compatible, that is, that the global incentive constraints are satisfied.¹⁷ We describe the numerical verification procedure in Appendix OA.2 of the Supplemental Material. For the range of parameters we study in Section 5, the allocations found using the second-order approach do indeed satisfy the global incentive constraints. In addition, the optimal allocations in such dynamic models with spillovers cannot easily—at this level of generality—be shown to be unique. However, we can show uniqueness for the functional forms and parameter values used in our simulations in Section 5.

¹⁷See also Farhi and Werning (2013) and Stantcheva (2017).

3.3. *Characterizing the Constrained Efficient Allocation Using Wedges*

To characterize the constrained efficient allocations, it is very helpful to define the so-called wedges or implicit taxes and subsidies that apply at these allocations. The wedges measure the distortions at the optimum relative to the laissez-faire economy with a patent system, that is, the hypothetical incentives expressed as implicit taxes or subsidies that would have to be provided to firms starting from the laissez-faire case in order to reach the allocation under consideration. The R&D effort wedge $\tau(\theta^t)$ measures the distortion on the firm’s R&D effort margin at history θ^t . It is equal to the gap between the expected stream of marginal benefits from effort and its marginal cost, where the expectation is conditional on the history θ^t . A positive wedge means that the firm’s effort is distorted downwards. This wedge will interchangeably be called the corporate tax or the profit wedge, since it will mimic a tax on firms’ profits, gross of R&D investments. The R&D investment wedge, or R&D wedge for short, $s(\theta^t)$ is defined as the gap between the marginal cost of R&D and the expected stream of benefits. It is akin to an implicit subsidy: a positive R&D wedge will mean that, conditional on the effort, the firm is encouraged to invest more in R&D than under laissez-faire with patent protection.

DEFINITION 1—The corporate wedge and the R&D wedge: The corporate (or profit) wedge is defined as

$$\tau(\theta^t) := \mathbb{E} \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \frac{\partial \lambda_t(\theta^t)}{\partial l_t(\theta^t)} \right) - \phi'(l_t(\theta^t)) \right).$$

The R&D spending (or R&D) wedge is defined as

$$s(\theta^t) := M'_t(r_t(\theta^t)) - \frac{1}{R} \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \frac{\partial \lambda_{t+1}(\theta^{t+1})}{\partial r_t(\theta^t)} \right) \right).$$

To simplify the notation, we use the following definitions:

$$\Pi_t(\theta^t) := \sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \pi_s(q(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)}$$

is the marginal impact on future expected profit flows from an increase in quality q_t . Let

$$Q_t(\theta^t) := \sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \tilde{Y}(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)}$$

be the marginal impact of quality on future expected output net of production costs, \tilde{Y} .

$$Q_t^*(\theta^t) := \sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \tilde{Y}^*(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)}$$

is the marginal impact on future expected output net of production costs from an increase in quality q_t , when quantity is set by the Planner to the socially optimal level.¹⁸

¹⁸Note that since the quantity maximizes consumption net of production costs per producer, that is, reaches $\tilde{Y}^*(q(\theta^s), \bar{q}_s)$, the derivative is just the direct impact of quality (the indirect effect through a change in the quantity is zero).

4. OPTIMAL POLICIES

In this section, we characterize the optimal constrained efficient allocations that are the solutions to the planning problem in Section 3. We then show how these allocations can be implemented with a parsimonious tax function.

4.1. *Optimal Corporate and R&D Wedges*

Denote by $\varepsilon_{xy,t}$ the elasticity of variable x to variable y at time t :

$$\varepsilon_{xy,t} := \frac{\partial x_t}{\partial y_t} \frac{y_t}{x_t}.$$

For instance, $\varepsilon_{l(1-\tau),t}$ is the elasticity of R&D effort to the net-of-tax rate $1 - \tau$. Denote by λ_x the partial derivative of the step size λ with respect to variable x . Taking the first-order conditions of program $P(\bar{q})$, and rearranging, yields the optimal wedge formulas at given \bar{q} in parts (i) and (ii) in the next proposition. Solving the full program yields an expression for the multipliers on the consistency constraints in part (iii), and hence a solution for \bar{q}_t .

PROPOSITION 1—Optimal corporate wedge and R&D wedge: (i) *The optimal profit wedge satisfies*

$$\begin{aligned} \tau(\theta^t) = & \underbrace{-\mathbb{E}\left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R}\right)^{s-t} \eta_s\right) \frac{\partial \lambda_t}{\partial l_t}}_{\text{Pigouvian correction}} - \underbrace{\mathbb{E}(Q_t(\theta^t) - \Pi_t(\theta^t)) \frac{\partial \lambda_t}{\partial l_t}}_{\text{Monopoly quality valuation correction}} \\ & + \underbrace{\frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t}(\theta^t)}_{\text{Type distribution and persistence}} \underbrace{\frac{\phi'_t \lambda_{\theta t}}{\lambda_t} \left[\frac{1}{\varepsilon_{l,1-\tau}} \frac{1}{\varepsilon_{\lambda l,t}} + \rho_{\theta l,t} \right]}_{\text{Elasticity}}. \end{aligned} \tag{9}$$

(ii) *The optimal R&D subsidy is given by*

$$\begin{aligned} s(\theta^t) = & \underbrace{\frac{1}{R} \mathbb{E}\left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R}\right)^{s-t-1} \eta_s \frac{\partial \lambda(\theta^{t+1})}{\partial r_t}\right)}_{\text{Pigouvian correction}} \\ & + \underbrace{\frac{1}{R} \mathbb{E}\left((Q_{t+1}(\theta^{t+1}) - \Pi_{t+1}(\theta^{t+1})) \frac{\partial \lambda(\theta^{t+1})}{\partial r_t}\right)}_{\text{Monopoly quality valuation correction}} \\ & + \underbrace{\frac{1}{R} \mathbb{E}\left(\frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t+1}(\theta^{t+1}) \phi'_{t+1}(l(\theta^{t+1})) \frac{\lambda_{\theta} \lambda_r}{\lambda \lambda_l} (\rho_{lr} - \rho_{\theta r})\right)}_{\text{Type distribution and persistence}} \underbrace{\left(\rho_{lr} - \rho_{\theta r}\right)}_{\text{Relative complementarity}}. \end{aligned} \tag{10}$$

(iii) *The multipliers η_t , capturing the spillovers between firms are given by*

$$\int_{\theta^t} \frac{\partial \tilde{Y}(q_t(\theta^t), \bar{q}_t)}{\partial \bar{q}_t} P(\theta^t) d\theta^t = \eta_t.$$

PROOF: See Appendix OA.1 of the Supplemental Material.

Q.E.D.

The optimal wedges in (9) and (10) are determined by the trade-off between maximizing allocative efficiency and minimizing informational rents. They balance three main effects:

(1) Monopoly quality valuation correction. The intermediate goods monopolist takes into account the effect of a quality increase on profits, while the planner values the effect on household consumption. Recall that the wedge is defined as the implicit subsidy (or implicit tax) starting from the laissez-faire allocation with patent protection. To induce the monopolist to invest more in quality than they would if they were maximizing profits, this term decreases the profit wedge and increases the R&D wedge. When quantity is chosen by the intermediate goods producer in the private market to maximize profits and not social surplus, the effect of a change in quantity (induced by extra R&D investment or R&D effort) on social welfare (implicit in $Q_t(\theta^t)$) is first-order and is proportional to the monopoly distortion, that is, the gap between price and marginal cost, summed over all future periods.¹⁹ This monopoly quantity correction term is positive and always makes the profit wedge smaller and R&D subsidy larger relative to a case where there is no difference between social and private valuation (i.e., no monopoly distortion and the producer perfectly internalizes social value on the production side). This is intuitive: the larger the gap between the monopolist’s value and the social value, the less the monopolist internalizes the social benefit from an increase in quality, and the more they need to be incentivized to invest in innovation.

Let us think of two polar cases. If there is no IPR at all in the laissez-faire economy, profits are zero and a large subsidy is needed to incentivize innovation. If, on the other hand, the laissez-faire features a prize system in which the company is entirely paid for the social value it generates, the monopoly distortion is zero and a smaller subsidy is needed to incentivize the investment in innovation. In between these polar cases, if profits are a share α_π of total net social value, the remaining gap in value that needs to be incentivized is $\mathbb{E}((1 - \alpha_\pi)Q_{t+1}(\theta^{t+1}))$; α_π can capture in a reduced-form way the share of total value granted to innovating firms by the IPR system. The closer the company is to capturing the full social surplus, the less additional incentive provision is needed.

(2) Pigouvian correction for the technology spillover. As long as the technological spillover is positive, the Pigouvian correction term unambiguously increases firms’ R&D effort and investment relative to laissez-faire. The Pigouvian correction for R&D effort in (9) is increasing in the effect of effort on the step size ($\frac{\partial \lambda_t}{\partial l_t}$). The correction for R&D spending in (10) is increasing in the expected effect of R&D investments on the next period’s step size $\frac{\partial \lambda_{t+1}}{\partial r_t}$.

¹⁹Formally, $\frac{\partial \tilde{Y}(q_t(\theta^t), \bar{q}_t)}{\partial q_t} = \frac{\partial Y(q_t(\theta^t), k_t(q_t(\theta^t), \bar{q}_t))}{\partial q_t(\theta^t)} + (p(q_t(\theta^t), k_t(q_t(\theta^t), \bar{q}_t)) - \frac{\partial C(k_t(q_t(\theta^t), \bar{q}_t))}{\partial k}) \frac{\partial k_t(q_t(\theta^t), \bar{q}_t)}{\partial q_t(\theta^t)})$, where $k_t(q_t(\theta^t), \bar{q}_t)$ is the quantity chosen to maximize profits by a monopolist with quality $q_t(\theta^t)$. This derivative also appears in the Pigouvian correction term: when aggregate quality increases, quantity produced increases, which has a first-order positive effect on social welfare.

While the monopoly distortion captures the lack of alignment in the valuation of quantity produced, the Pigouvian correction captures the lack of alignment on how much quality produced is valued socially and privately. Even if there is no monopoly power at all, this distortion applies.

(3) Screening and incentives. Screening considerations may push in the opposite direction from the monopoly and Pigouvian corrections. The screening term arises because of asymmetric information. Without asymmetric information, this term would be zero and the optimal profit wedge and the optimal R&D subsidy would be equal to the Pigouvian and monopoly quality valuation corrections, as in Section 2.3. Externalities would be corrected under full information (and tailored to each research productivity history θ'), and there would be no informational rents. With asymmetric information, there are three effects at play.

The stochastic process for firm type. The initial type distribution times the persistence in types (captured by the impulse response function $I_{1,t}$) increases the magnitude of the profit wedge and the R&D investment wedge. More persistent types effectively confer more private information to firms and, hence, higher potential informational rents. To reduce these informational rents, allocations have to be distorted more (the typical trade-off between informational rents and efficiency). If shocks were i.i.d., we would have $I_{1,t} = 0$ for all $t > 1$, and, hence, the optimal corporate and R&D wedges would be equal to only the Pigouvian correction term plus the monopoly valuation correction term for all $t > 1$. With AR(1) shocks with persistence parameter \tilde{p} , $I_{1,t} = \tilde{p}^{t-1}$ so that the impulse response is fully determined by the persistence parameter. If types are fully persistent, so that there is only heterogeneity, but no uncertainty, the impulse response $I_{1,t} = 1$ for all t and the screening term does not decay over time.

A higher inverse hazard ratio $\frac{1-F^1(\theta_1)}{f^1(\theta_1)\theta_1}$ (implying a larger mass of firms with research productivity larger than θ_1 relative to the mass of firms with type θ_1 ($f^1(\theta_1)$)) makes the cost of inducing a marginal distortion in effort or R&D investments at point θ_1 small relative to the benefit of saving on the informational rent over a mass of $1 - F^1(\theta_1)$ of all firms more productive than θ_1 .

The efficiency cost of distorting R&D effort. A higher efficiency cost decreases the optimal effort wedge.²⁰ The efficiency cost can be decomposed into allocative inefficiency and information rents. The allocative inefficiency induced by the effort wedge is increasing in the elasticity of the step size with respect to effort ($\varepsilon_{l,1-\tau,t}\varepsilon_{\lambda l,t}$). The informational rent inefficiency increases in the complementarity of effort to firm research productivity $\rho_{\theta l,t}$. A high complementarity between effort and firm type means that it is easy for higher research productivity firms to mimic lower productivity ones, which increases their potential informational rent, and thus leads to an optimally higher distortion in the allocation to reduce these rents. Since the disutility of R&D effort is indexed by t , the strength of this incentive effect could vary over the life cycle of a firm.

The complementarity between R&D, firm effort, and firm type. For the purposes of screening, observable R&D investments are distorted only in so far as they can indirectly affect the unobservable R&D effort choice, that is, can affect the incentive constraint of the high research productivity firm.

How effective R&D investment subsidies are at stimulating unobserved effort depends on the relative complementarity of R&D expenses with effort and type, ($\rho_{lr} - \rho_{\theta r}$), which determines the sign of the screening term. Higher R&D expenses lead to more effort by

²⁰This is naturally reminiscent of the inverse elasticity rule in Ramsey taxation.

the firm as long as they increase the marginal return to effort, that is, as long as $\frac{\partial^2 \lambda(l, r, \theta)}{\partial r \partial l} > 0$ and thus $\rho_{lr} > 0$, as seems likely. On the other hand, if $\rho_{\theta r} > 0$, then higher R&D expenses have a higher marginal effect on the step sizes of high research productivity firms (at any given effort level), which makes it easier for them to mimic the step sizes allocated to lower productivity firms. This, in turn, increases the informational rent that needs to be given to these firms to induce them to reveal their true type. What matters is whether, on balance, the net effect of increasing R&D is positive, that is, whether the effect on effort will outweigh the effect on the step size conditional on effort. If yes, then R&D expenses will relax the firms' incentive constraints and reduce their informational rents. This occurs if $(\rho_{lr} - \rho_{\theta r}) > 0$, that is, if R&D expenses are more complementary to effort than they are to firm type.

If the complementarity of R&D with both R&D effort and firm type is the same ($\rho_{lr} = \rho_{\theta r}$), then the screening term of the optimal R&D subsidy is zero. In this special case, an increase in R&D has exactly offsetting effects on effort and on the step size conditional on effort, leaving the informational rents unchanged on balance (i.e., the incentive constraints are unaffected by changes in R&D investments).

Another way of interpreting $\rho_{\theta r}$ is as the riskiness of R&D, or as its exposure to the intrinsic risk of the firm. The higher this complementarity, the more R&D returns are subject to the stochastic realizations of firm type. Hence, the sign of $(\rho_{lr} - \rho_{\theta r})$ measures the strength of R&D contribution to firm effort, filtered out of the exposure to firm risk.

In general, there is no reason to think that the Hicksian coefficients of complementarity are constant. They could vary with the level of effort, R&D, and ability, as well as with firm age.²¹ Hence, the optimal R&D wedge may change sign over the distribution of types or over the life cycle of a firm.

4.1.1. *Cross-Sectional Profile of Optimal Policies*

At this level of generality, we cannot pin down how the optimal wedges vary with firm productivity. However, we can discuss what forces drive each term and the cross-sectional patterns. It is important to bear in mind that a higher R&D wedge does not mean a higher investment in R&D, and a lower effort wedge does not mean more R&D effort. It merely means a higher incentive relative to laissez-faire. This is because firms have heterogeneous benefits and costs from investments and effort under laissez-faire, so that the same level of incentive will not translate into the same level of inputs across firms. For instance, in the laissez-faire, low research productivity firms invest much less than high research productivity firms and this pattern is not overturned despite the incentive provision.

The screening term will tend to be larger in absolute value for lower productivity firms, by the logic of screening models: because higher productivity firms want to pretend to be lower types, lower types' allocations will be distorted to prevent such deviations while also minimizing informational rents. For the profit wedge, the screening term is positive, which means that higher type firms will face a lower profit wedge. For the R&D wedge, the screening term's sign depends on $\rho_{rl} - \rho_{\theta r}$. When $\rho_{\theta r} > \rho_{rl}$, the screening term is negative. R&D investments disproportionately benefit high productivity firms. It is then better not to incentivize R&D investments as much for lower productivity firms, since this makes their allocations more attractive to high productivity firms. Since in this case

²¹Although we have dropped this notation for clarity, all elasticities, coefficients of complementarities, and functions are evaluated at θ^t , so they can depend on investment size and on age t .

lower productivity firms have no comparative advantage at innovation, they should not be incentivized as much to invest in R&D, so that high productivity firms can be incentivized more.

Regarding the Pigouvian correction term, higher research productivity firms have a higher positive spillover on other firms as long as $\rho_{lr} > 0$ and $\rho_{\theta l} > 0$, in which case their marginal investment in R&D or a higher effort has a higher marginal impact on their step size, and hence on aggregate quality. The optimal Pigouvian correction would then be increasing in firm type and warrant a higher profit subsidy, all else equal, for higher productivity firms. The comparative statics of the monopoly valuation correction term are ambiguous. If private valuation is a constant share of total social valuation, then the monopoly valuation term would be increasing in firm type as well.

In addition, since high productivity firms invest more in R&D and generate larger profits, the statements just made about wedges will be the same if expressed in terms of observables, namely, profits and R&D investments.

4.1.2. Age Profile of Optimal Policies

The optimal policies will generically change over firms' life cycles. The first reason why policies depend on firm age is that they are set at age 1 under full commitment of the planner. As a result, it is the time that has passed relative to that initial period that induces age patterns. The screening term in the optimal corporate and R&D wedges declines in absolute value with age, as long as the impulse response is below 1 (as is the case with a first-order autoregressive or geometric autoregressive process with persistence parameter $\tilde{\rho} < 1$). This decay towards zero is faster the lower the persistence in types. From the perspective of period 1, since types are stochastic, the informational rent to be received after any particular history θ^t a longer time span away is worth less to the agent and is less costly to the planner. The smaller effective informational rents warrant less distortion in the allocations.

Hence, over the life cycle of a firm, the wedges converge to the Pigouvian and monopoly correction terms. Whether they converge from above or below depends on the sign of the screening term, which depends on the relative Hicksian complementarities of R&D investments to R&D effort versus unobserved productivity. If $\rho_{\theta r} > \rho_{lr}$, the screening term is negative and optimal wedges converge from below to the Pigouvian and monopoly correction terms. They converge from above if $\rho_{\theta r} < \rho_{lr}$.

The second reason why policies depend on age is that the technological fundamentals, such as the step size λ_t , the cost of effort $\phi_t(l_t)$, the distribution of types, and the cost of R&D $M_t(r_t)$, can vary with age. For instance, as a firm gains expertise with age, the cost of unobservable and observable R&D inputs may decrease. More empirical work could shed light on the life-cycle patterns of the production and innovation technologies. The age patterns of optimal policies are thus theoretically ambiguous and will depend on the parameters of the model.

4.2. Implementation Through Taxes and Subsidies

In this section, we show that the optimal allocations can be implemented with a relatively parsimonious tax function.²² In general, the optimal policies depend on the histories of R&D inputs and outputs in a nonlinear and non-separable way. However, there exists

²²Until now, we have considered a direct revelation mechanism, in which firms report their type to the planner every period and the planner assigns allocations as a function of the history of reports received. We

a simpler implementation of the optimal mechanism, which does not depend on histories longer than two periods. This makes our implementation very different from dynamic income tax models for agents as in [Farhi and Werning \(2013\)](#) or [Stantcheva \(2017\)](#), where it is in general impossible to cut the history-dependence of optimal policies except in special cases such as i.i.d. shocks ([Albanesi and Sleet \(2006\)](#)).

Market Structure. The constrained efficient allocations solved for in Section 3 are independent of the underlying market structure as long as the information set and toolbox of the planner are as specified there.²³ However, the shape and level of the tax function that implements the constrained efficient allocation depends on the market structure. For instance, the more credit constrained firms are in the laissez-faire decentralized market, the more generous transfers they would have to receive early on so as to be able to invest the amount required in the constrained efficient allocation. The implementation also depends on the IPR policy, which determines the level of profits under laissez-faire. Finally, the same optimal allocations can often be implemented with multiple different policies and, hence, the implementation is not unique.

We assume that in the laissez-faire market firms can borrow freely at a constant rate R , and that they take the price of the final good (normalized to 1) as given. They face the demand function for their differentiated intermediate goods under a patent system that grants full monopoly power, as presented in Section 2.

Implementation Result. The tax implementation function can be relatively parsimonious when the impulse response functions $I_{1,t}(\theta^t)$ are independent of the history of types, except through θ_1 and θ_t for all t , as would be the case for any AR(1) process, or a geometric random walk (or, for any monotonic transformation of an AR(1) process).

The constrained efficient allocation from program $P(\bar{q})$ is implemented with a comprehensive, age-dependent tax function that conditions on current quality q_t , lagged quality q_{t-1} , current R&D r_t , lagged R&D r_{t-1} , and first-period quality q_1 , that is, $T_t(q_t, r_t, q_{t-1}, r_{t-1}, q_1)$. The proof is in the Supplemental Material. Note that because profits are an immediate function of quality q_t , the tax function could instead be written as a function of profits, that is, $\tilde{T}_t(\pi_t, r_t, \pi_{t-1}, r_{t-1}, \pi_1)$.

It may at first glance seem as if this result were trivial: if productivity follows a Markov process, it appears to make sense that one only needs to condition on allocations one period back in addition to the current allocations. However, this intuition is not correct. Even with Markov shocks, most dynamic tax problems ([Farhi and Werning \(2013\)](#), [Stantcheva \(2017\)](#)) require conditioning on full histories. What is different here is that the past stock of quality can serve as a sufficient statistic for past investments, in the sense that it fully determines the future benefit from more innovation investments (together with r_t and r_{t-1}).

would now like to step away from reporting types and move into the realm of policy implementation. The question of implementation is whether there is some general tax and transfer function $T(q^\infty, r^\infty)$ that depends on the full sequence of all observables, that is, on the history of quality q^∞ (or, interchangeably, step size λ^∞) and R&D investment r^∞ , such that, if this tax and transfer rule is in place, optimizing firms will pick allocations equal to the constrained efficient allocation from the direct revelation mechanism.

²³For instance, if firms are credit constrained, the planner will simply increase the transfer in a lump-sum fashion in earlier periods and make up for it with lower transfers in later periods without affecting the incentive constraints. However, if the information set of the planner is altered, for example, if firms could save in a hidden way, then the constrained efficient allocation would be different.

The link between the wedges and the tax function is as follows, where the arguments of the tax function are evaluated at their optimal values for history θ^t :

$$s(\theta^t) = -\frac{1}{R} \mathbb{E} \left[\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \left(\frac{\partial T_s}{\partial q_s} + \frac{1}{R} \frac{\partial T_{s+1}}{\partial q_s} \right) \frac{\partial \lambda_{t+1}}{\partial r_t} \middle| \theta_t \right] \\ - \mathbb{E} \left[\sum_{s=t}^{\infty} \left(\frac{1}{R} \right)^{s-t} \left(\frac{\partial T_s}{\partial r_s} + \frac{1}{R} \frac{\partial T_{s+1}}{\partial r_s} \right) \middle| \theta_t \right], \\ \tau(\theta^t) = \mathbb{E} \left[\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \left(\frac{\partial T_s}{\partial q_s} + \frac{1}{R} \frac{\partial T_{s+1}}{\partial q_s} \right) \frac{\partial \lambda_t}{\partial l_t} \middle| \theta_t \right].$$

Generically, the optimal wedges in Proposition 1 depend nonlinearly and non-separably on the choice variables. Few general statements can be made without specifying the underlying functional forms. However, it is clear from the optimal wedge formulas that the degree of nonlinearity of the profit tax and R&D subsidy depends crucially on the shape of the step size function. In particular, profit taxes will be far from linear if the marginal effect of R&D effort is very nonlinear. If the step size is close to linear in R&D effort, on the contrary, then profit taxes will also be closer to linear. The same applies to the R&D wedge with respect to the marginal effect of R&D investment on the step size. Similarly, nonseparabilities in the tax function between profits and R&D investment will be quantitatively important if the step size features strong complementarities between R&D effort and investment.

In Appendix OA.2 of the Supplemental Material, we work out a very simple functional form example that generates constant marginal profit taxes and constant R&D subsidies, which only depend on the degree of market power (β) and the strength of the spillover (ζ).

5. QUANTITATIVE INVESTIGATION

In this section and the next one, we provide empirical content to the theoretical model by estimating it and numerically illustrating the optimal policies. We present and discuss estimated values of the key parameters, such as the complementarity between R&D and firm research productivity ρ_{θ} , the persistence \bar{p} , and the externality strength ζ . We also document the age and cross-sectional patterns of optimal policies and allocations. Finally, we assess the losses from simpler policies.

5.1. Data and Summary Statistics

The theory developed in this paper can be applied to different data sets, and our model could be estimated for different countries, industries, and types of firms to inform the specific optimal policies for each setting or sample under consideration. The benchmark data we chose is the Census Bureau's Longitudinal Business Database and Census of Manufacturers, matched to the U.S. Patent and Trademark Office (USPTO) patent data from the NBER database (as described in detail in Hall, Jaffe, and Trajtenberg (2001)), containing over three million patents with their forward citations. This data match was done and used by Acemoglu et al. (2018) and Akcigit and Kerr (2018). We externally calibrate some parameters of our model using estimates provided in the literature, as described below. We then internally estimate the remaining parameters.

An alternative data set is the firm-level accounting data from COMPUSTAT matched to the NBER patent database. As argued in Bloom, Schankerman, and Van Reenen (2013), these firms represent a large fraction of the innovation in the United States. For completeness, we provide numerical results based on the COMPUSTAT sample in Section S.3 within the replication file.

Map Between the Model and the Data. One advantage of the patent data matched to firm-level data is that there is a natural mapping between the variables in the model and the data. R&D spending $M(r)$ can directly be measured as reported R&D expenses in the accounting data. The step size λ_t , that is, the flow of new quality of a firm in year t , can be measured by the forward citations received on all innovations patented in year t . The quality q_t is the depreciation-adjusted stock of citations per patent to date, $q_t = (1 - \delta)q_{t-1} + \lambda_t$. We can also directly measure profits and sales.

5.2. Estimation

We first parameterize the model as summarized in Table I. Some of the parameters are calibrated exogenously, following the earlier innovation literature. This reduces the size of the parameter vector to be estimated. We report these parameters in the upper panel of Table II. In Section 5.5, we provide sensitivity analyses for each parameter. The lower panel of Table II reports the key parameters of our model, which are estimated internally to best match important moments in the data presented in Table III. In Table IV, we check that the estimation does a good job matching non-targeted moments to assess the fit. We describe our estimation procedure in more detail now.

The Status Quo Economy. To be able to consistently estimate the parameters of the model by matching moments in the data, we need to subject firms in our model to the same policies (R&D subsidy and corporate tax) as in the United States. We call our baseline setting the *status quo* economy. This setting has the same primitives as the setting considered until now, but instead of optimal policies it features the current policies in place in the U.S. We approximate real-world R&D subsidies with a linear R&D subsidy rate. We estimate the effective subsidy rate on R&D investments by firms using the total

TABLE I
FUNCTIONAL FORMS.

Function	Notation	Functional Form
Consumer valuation	$Y(q_t, k_t)$	$\frac{1}{1-\beta} q_t^\beta k_t^{1-\beta}$
Cost function	$C_t(k, \bar{q}_t)$	$\frac{k}{\bar{q}_t^\gamma}$
Quality accumulation	$H(q_{t-1}, \lambda_t)$	$q_t = (1 - \delta)q_{t-1} + \lambda_t$
Step size	$\lambda_t(r_{t-1}, l_t, \theta_t)$	$(\alpha r_{t-1}^{1-\rho_{\theta r}} + (1 - \alpha)\theta_t^{1-\rho_{\theta r}})^{\frac{1}{1-\rho_{\theta r}}} l_t$
Disutility of effort	$\phi_t(l_t)$	$\kappa l_t^{\frac{1+\gamma}{1+\eta}}$
Cost of R&D	$M_t(r_t)$	$\kappa_r r_t^{\frac{1+\eta}{1+\gamma}}$
Stochastic type process	$f^r(\theta_t \theta_{t-1})$	$\log \theta_t = \bar{p} \log \theta_{t-1} + (1 - \bar{p})\mu_\theta + \epsilon_t$
Distribution of heterogeneity θ_1	$f^1(\theta_1)$	$f^1(\theta_1) = \frac{I_{\Theta_1}(\theta_1)}{\theta_1[\underline{\theta}_1 - \bar{\theta}_1]}$
Initial quality level	q_0	0

Note: $I_{\Theta_1}(\theta_1)$ denotes the indicator function equal to 1 if θ_1 is in the set $\Theta_1 = [\underline{\theta}_1, \bar{\theta}_1]$.

TABLE II
PARAMETER VALUES.

Parameter	Symbol	Value	Standard Error
<i>External Calibration</i>			
Interest rate	R	1.05	
Intangibles depreciation	δ	0.1	
Knowledge share	β	0.15	
R&D cost elasticity	η	1.5	
Level of types	μ_θ	0.00	
Initial R&D stock	r_0	1.0	
Program horizon	T	30	
<i>Internal Calibration</i>			
R&D share	α	0.483	(0.025)
R&D-type substitution	$\rho_{\theta r}$	1.88	(0.126)
Type variance	σ_ϵ	0.320	(0.014)
Type persistence	\bar{p}	0.63	(0.022)
Scale of disutility	κ_l	0.69	(0.050)
Scale of R&D cost	κ_r	0.055	(0.003)
Effort cost elasticity	γ	0.86	(0.052)
Support width for θ_i	Θ^1	1.91	(0.097)
Production externality	ζ	0.018	(0.001)

spending of the government on firm R&D through all programs (R&D tax credits, direct grants, etc.) divided by total private business spending on R&D. This yields an average effective subsidy rate of 19%. The details for this computation are in the Supplemental Material. The average effective corporate tax rate is set at 23%. Reassuringly, the estimation procedure for our key parameters is not very sensitive to the choice of these effective rates.

Functional Forms. The cost function depends on aggregate quality \bar{q}_t , and the strength of the externality is measured by ζ . The step size is multiplicatively separable in R&D effort l_t and takes a constant elasticity of substitution (CES) form in type θ_t and R&D investment r_{t-1} . In this case, $\rho_{\theta l} = \rho_{lr} = 1$. We specify this functional form for tractability; given the data, it would be very difficult to empirically discipline ρ_{lr} . Given that the sign

TABLE III
MOMENTS.

Moment	Target	Simulation	Standard Error
M1. Patent quality-R&D elasticity	0.88	0.97	(0.0009)
M2. R&D/Sales mean	0.041	0.035	(0.0025)
M3. Sales growth (DHS) mean	0.06	0.07	(0.005)
M4. Within-firm patent quality coeff of var	0.63	0.76	(0.0017)
Across-firm patent quality coeff of var:			
M5. Young firms	1.06	1.05	(0.0012)
M6. Older firms	0.99	0.81	(0.0016)
M7. Patent quality young/old	1.04	1.08	(0.0048)
M8. Spillover coefficient	0.191	0.192	(0.046)
M9. Elasticity of R&D investment to cost	-0.35	-0.35	(0.101)

TABLE IV
GOODNESS OF FIT FOR NON-TARGETED MOMENTS.

	Data	Model
Sales growth bottom 90% vs. top 10%	0.03	0.04
R&D/Sales for bottom 90%	0.038	0.034
R&D/Sales for top 10%	0.052	0.042
Ratio sales for old firms vs. young firms	1.73	2.32

of $\rho_{lr} - \rho_{or}$ determines the sign of the screening term in the optimal R&D subsidy (as shown in Proposition 1), the key question for whether screening will lead to a higher or lower subsidy on R&D will be whether $\rho_{or} \geq 1$ or $\rho_{or} < 1$. The costs of R&D effort and R&D investments are iso-elastic. Finally, the stochastic process for firm research productivity type is a geometric random walk, with persistence \tilde{p} (additional shock processes are in the replication file). The shock ε_t follows a normal distribution with mean zero and variance σ_ε .

Externally Calibrated Parameters. We take the externally calibrated parameters from reputable papers in the innovation and growth literatures, but we also provide many sensitivity analyses in Section 5.5.

The profit parameter β is set to 0.15 as in Guner, Ventura, and Xu (2008). The exponent on the R&D cost function, η , is set as in Akcigit and Kerr (2018). The depreciation parameter δ is a standard feature of empirical innovation work, as taken from Hall, Jaffe, and Trajtenberg (2005). This depreciation reflects the idea that, from the point of view of each individual innovation-producing firm, knowledge can become obsolete unless it is updated, which carries a cost. The long-run discount rate R reflects the interest rate plus the probability of exit (or death). The average level of research productivity is normalized to $\mu_\theta = 0$, while the initial R&D stock is set to $r_0 = 1$.

Moments and Identification. Table III lists the data moments that we match. The second column provides the value of the moment in the simulations, the third column gives the target value of each moment in the data, and the fourth column shows the standard error. In this section, we discuss the identification of the parameters in our model.

Let the vector of the nine endogenously estimated parameters be denoted by

$$\mathbb{X} = (\alpha, \rho_{or}, \sigma_\varepsilon, p, \kappa_l, \kappa_r, \gamma, \zeta, \Theta^1).$$

In our benchmark estimation, we chose the parameters to minimize the loss function:

$$L(\mathbb{X}) = \sum_{k=1}^9 \left(\frac{\text{moment}_k^{\text{model}}(\mathcal{X}) - \text{moment}_k^{\text{data}}}{\text{moment}_k^{\text{data}}} \right)^2,$$

where $\text{moment}_k^{\text{model}}$ is the value of moment k in the model and $\text{moment}_k^{\text{data}}$ is the value of the moment in the data. In the replication file, we estimate the parameters using a two-step GMM-type weighting.

Since we are minimizing the weighted distance between the theoretical and empirical moments, all parameters are identified jointly. Nevertheless, given the dynamics in our model, we can provide a heuristic discussion of identification.

Elasticity of Patent Quality wrt. R&D, M1: The first moment is the elasticity of patent quality with respect to R&D spending, where patent quality is measured as citations per patent. This moment measures how effective R&D spending is at generating successful innovations. It has been estimated in the literature since Griliches (1998). Not surprisingly, this moment informs the complementarity (or elasticity of substitution) parameter $\rho_{\theta r}$ in the innovation production function.

R&D Intensity, M2: The second moment is the mean ratio of R&D spending to firm sales, which is a measure of the R&D intensity of a firm. It was computed by Acemoglu et al. (2018) and is consistent with other papers. The R&D share in the step size, α , affects the marginal return to R&D investment r_t and therefore has a direct impact on firms' R&D/Sales ratio.

Sales Growth, M3: The third moment we include is firms' sales growth. Firm growth is determined by R&D investments. These are in turn driven by the firms' first-order condition that sets the marginal return from R&D investment equal to its marginal cost. Therefore, the scale parameter of the cost function, κ_r , has a first-order impact on the average growth rate of the firm.

Within-firm Patent Quality Variation, M4: The fourth to sixth moments are specific to our model, which highlights the role of firm heterogeneity and the role of uncertainty over time. Moment four considers the variation in a firm's quality (again, as measured by its citations per patent) over time. This within-firm measure helps assess the uncertainty facing a firm, which is captured by the persistence parameter \tilde{p} in our model.

Across-firm Patent Quality Variation by Age, M5–M6: The fifth and sixth moments capture the variation in quality across firms. This cross-sectional variability measure gauges the degree of heterogeneity across firms and is computed separately for young and old firms. "Young" firms are defined—both in the data and in the model—as those of age 0–5 years. "Old" firms are older than 5 years (we tried alternative definitions of young and old, with cutoffs at 3 or 10 years, with extremely similar results). These moments are mainly determined by the dispersion σ_ϵ and the width of the support of the type distribution Θ^1 .

Patent Quality Ratio (young/old), M7: The seventh moment is the ratio of patent quality between young and old firms and measures the decline in invention quality that occurs with firm age.

Spillover Coefficient, M8: One of the key moments, M8, targets the estimate of technological spillovers in Bloom, Schankerman, and Van Reenen (2013). These authors estimated spillovers by regressing the sales of a firm on the R&D of other firms in the economy, weighted by the extent of technological proximity with these other firms. They instrumented for this R&D using exogenous variation in effective R&D tax credit rates at the firm level. We estimate the spillover parameter ζ in our model through indirect inference. More precisely, we replicate their instrumental variable regression by exogenously shocking \bar{q}_t and generating simulated economies. We then regress the sales in the model on the R&D of other firms in the economy and match the regression coefficient to the one in Bloom, Schankerman, and Van Reenen (2013). We obtain a very close fit. This process helps us identify the spillover strength ζ .

Elasticity of R&D investment to R&D costs, M9: The final moment is the elasticity of R&D investments to R&D costs, taken from Bloom, Griffith, and Van Reenen (2002). They found an elasticity of R&D to user costs of -0.35 , which our model is able to match very closely.

Goodness of Fit: Non-targeted Moments. To check whether the fit of our estimated model is good even for non-targeted moments, we provide the values of four important

and non-targeted moments in the data and the model in Table IV, which pertain to the life cycle of firms or to the skewness and tails of the sales and R&D distributions. These are the sales growth of the bottom 90% firms versus the sales growth of the top 10% firms (as ranked by sales growth in their first 5 years); the ratio of sales for old versus young firms; and the R&D intensities (i.e., R&D divided by sales) for the bottom 90% versus top 10% firms. The fit is quite good, lending further credibility to our estimation and confirming that we are able to capture the more detailed tail behavior of the data.

5.3. Results

Table II shows the estimated parameters of the model. Focusing on some of the key parameters that were highlighted in Section 4, we see that, based on the data, R&D investments are highly complementary to firm research productivity: highly productive firms are disproportionately good at transforming R&D inputs into innovation. The type persistence is moderate, with $\tilde{p} = 0.63$. We can now simulate the optimal allocations and wedges, which we presented in analytical form in Section 4.

Gross Incentives and Net Incentives. A brief discussion of gross and net incentives for R&D is useful here (and in practice) when thinking about the magnitudes of incentives or disincentives actually provided for R&D. Let us illustrate the difference with linear taxes, to make the discussion simpler. If the profit tax applies to profits gross of R&D spending, that is, if R&D expenses are not deductible from the corporate tax base, the gross subsidy rate \tilde{s} is such that the firm's per-period payoff is

$$\pi(1 - \tau) - (1 - \tilde{s})M(r).$$

The net incentive on R&D—the rate that would apply to R&D expenses if they were also deductible from the profit tax base—is denoted by s and is defined such that the payoff of the firm is

$$\pi - \tau(\pi - M(r)) - (1 - s)M(r).$$

The net incentive is not well-measured by the gross subsidy rate \tilde{s} , since the profit tax captures part of the return to R&D investments. Thus, a share of the gross subsidy simply goes towards canceling out the disincentive effect from the profit tax. The net incentive is driven by the difference between the gross linear subsidy \tilde{s} and the tax τ : $s = \tilde{s} - \tau$. Put differently, there are two ways to incentivize R&D: either tax its returns less (the $-\tau$ term), or subsidize its costs more (the \tilde{s} term). As we will see, for screening purposes one way may be better than the other. Converting the combination of corporate income taxes and subsidies to a “net incentive” is also of great practical use, since different countries' systems load incentives on different parts of the tax code.

To highlight this distinction, in figures below, in addition to depicting $s(\theta')$, we also show the “gross” R&D wedge $\tilde{s}(\theta')$, namely, the gap between marginal costs and marginal benefits of R&D, taking into account the R&D effort wedge, that is, that there is simultaneously a tax on profits. Furthermore, to facilitate an interpretation of the wedges as tax and subsidy rates, we slightly redefine the profit and R&D wedges as fractions of profits and R&D costs. The R&D subsidy rate $\tilde{s}(\theta')$ is (locally) the fraction of the cost $M(r)$ that the firm does not have to pay, while the profit wedge $\tilde{\tau}(\theta')$ is the fraction of profits that

the firm pays:²⁴

$$\begin{aligned} \tilde{s}(\theta^t) &= 1 - \frac{1}{R} \frac{1}{M'_t(r(\theta^t))} \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} (1-\tilde{\tau}(\theta^t)) \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \frac{\partial \lambda_{t+1}(\theta^{t+1})}{\partial r_t(\theta^t)} \right), \\ & (1-\tilde{\tau}(\theta^t)) \mathbb{E} \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \right) \frac{\partial \lambda_t(\theta^t)}{\partial l_t(\theta^t)} = \phi'_t(l_t(\theta^t)). \end{aligned}$$

In the linear example above, s is directly comparable to the wedge $s(\theta^t)/M'(r)$ from Section 4, while \tilde{s} is comparable to $\tilde{s}(\theta^t)$.

5.3.1. *Cross-Sectional Patterns of the Optimal Allocations*

Panels C and D of Figure 2 plot the optimal profit wedge $\tilde{\tau}(\theta^t)$, the gross R&D wedge $\tilde{s}(\theta^t)$, and the corresponding net R&D wedge $s(\theta^t)/M'(r)$ for firms of different profit levels and R&D investments for ages $t = 2, 5, \text{ and } 15$. Panels E and F also depict these same wedges, but against unobservable productivity on the horizontal axis.

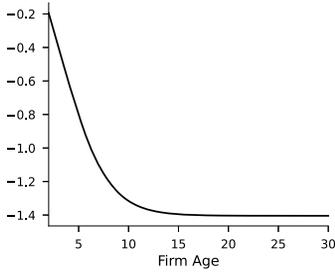
The wedges on profits are negative, while those on R&D investment are positive. This means that, on both the effort and R&D investment side, firms are incentivized to provide more of these inputs than they would under *laissez-faire*. Thinking back to the optimal wedge formulas in Section 4, this is to account for the monopoly distortion effect (incentivize monopolists to produce more indirectly) and the Pigouvian correction effect (to correct for the spillover), while still screening firms. It is worth clarifying that these are of course akin to marginal taxes (or subsidies), not average or total taxes. On balance, the government is still raising positive net revenues and consumers still get to consume a positive net output.

Let us consider how wedges vary by firm type, profit levels, and R&D expenses, remembering the theoretical discussion in Section 4. Incentives are described by the profit wedge and the net R&D wedge $s(\theta^t)$. For any given levels of the monopoly distortion and the Pigouvian correction terms, the screening term is larger in absolute value for lower type firms. This is the logic of screening models: since higher type firms are tempted to pretend to be lower types, lower types firms' allocations are distorted to prevent higher types from lying. For the profit wedge, the screening term is positive, which means that higher productivity firms will face a lower profit wedge (i.e., a less positive marginal profit tax or a higher marginal profit subsidy). For the R&D net wedge, the screening term's sign depends on $1 - \rho_{\theta r}$. According to our estimation, $\rho_{\theta r} > 1$, so the screening term is negative. Hence, the net R&D wedge is larger for higher type firms. When $\rho_{\theta r} > 1$, R&D investments disproportionately benefit high productivity firms. It is better to incentivize R&D investments less for the lower productivity firms, as this makes mimicking more attractive

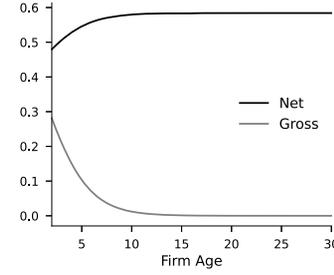
²⁴ $\tilde{s}(\theta^t)$ and $\tilde{\tau}$ are related to the wedges from Definition 1 through

$$\begin{aligned} \tilde{s}(\theta^t) &= s(\theta^t) \frac{1}{M'_t(r(\theta^t))} + \tilde{\tau}(\theta^t) \frac{1}{R} \frac{1}{M'_t(r(\theta^t))} \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \frac{\partial \lambda_{t+1}(\theta^{t+1})}{\partial r_t(\theta^t)} \right) \\ \text{and } \tilde{\tau}(\theta^t) &= \frac{\tau(\theta^t)}{\mathbb{E} \left(\Pi_t(\theta^t) \frac{\partial \lambda_t(\theta^t)}{\partial l_t(\theta^t)} \right)}. \end{aligned}$$

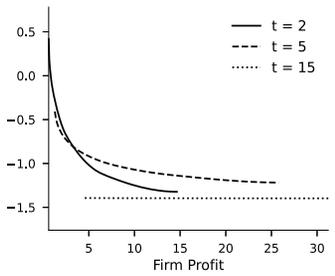
(a) Profit Wedge by Age



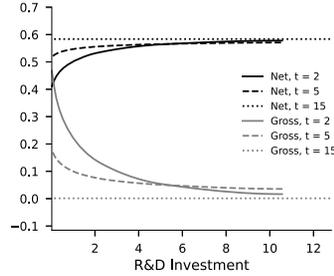
(b) R&D Wedges by Age



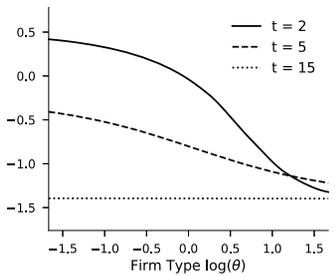
(c) Profit Wedge as Function of Profits



(d) R&D Wedges as Functions of R&D Investments



(e) Profit Wedge as Function of Type θ_t



(f) R&D Wedges as Functions of Type θ_t

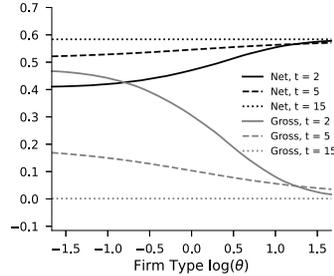


FIGURE 2.—Optimal profit and R&D wedges.

for high productivity firms.²⁵ In short, lower productivity firms that have no comparative advantage at innovation are not incentivized as much to invest in R&D, so that high productivity firms can be incentivized more. Naturally, they are still incentivized to some (possibly even to a large) extent because of the monopoly and Pigouvian corrections.

This logic is illustrated in Panels E and F, where the screening terms of the very high type firms converge almost entirely to the monopoly and the Pigouvian correction terms. Since high type firms also invest more in R&D and have larger profits, the wedges follow the exact same pattern when plotted against observables, namely, profits (panel C) and R&D investments (panel D). Recall that the net wedge summarizes the incentive provi-

²⁵Recall that a higher R&D wedge does not mean a higher investment in R&D; it just means a higher incentive relative to the laissez-faire. Under laissez-faire, low research productivity firms already invest much less than high research productivity firms and this pattern is not overturned despite the incentive provision.

sion for innovation. However, the gross wedge more closely matches the intuitions that come from explicit subsidies or taxes. The gross wedge is smaller for higher type firms, because it is partially compensating for the (lower) profit wedge, but the net incentive for R&D provided is larger for higher productivity firms.

High productivity firms are *on net* more incentivized to invest in innovation, and this incentive comes from a lower profit wedge rather than from a higher gross R&D wedge. This is the best mechanism for screening: higher productivity firms will be able to generate more profits from the same research investments, so the way to attract them to a given allocation that features more R&D investments (without attracting low productivity firms) is by letting the profit wedge at that allocation be lower, rather than by making the R&D wedge higher. This will encourage high productivity firms to put in more of the unobserved innovation input, which cannot be directly subsidized. In some sense, this is “performance-based” taxation, where good performance, rather than simply more (observable) inputs, is rewarded.

Keep in mind, however, that this discussion is still about wedges. When it comes to the (approximate) implementation using simpler policies that we consider in Section 6, a low profit wedge for higher research productivity firms can *approximately* be achieved in several ways. The most immediate policy that perfectly mirrors the wedges features lower marginal profit taxes on more profitable firms and a lower marginal subsidy at higher levels of R&D investments. But other implementations may work almost as well, if the loss from fine-tuning is small. For instance, a constant profit tax that is more generous than it should be for low profit firms, and at about the right level for high profit firms, could do reasonably well if the loss from giving low profit firms an excessively generous tax is quantitatively small.

Panels C and D of Figure 3 show the optimal inputs for firms of different productivities for different ages. Higher research productivity firms should optimally provide more effort and invest more in R&D. Given that the estimated parameters imply that $\rho_{l0} > 0$ and $\rho_{r0} > 0$, effort and R&D expenses for higher productivity firms have higher marginal benefits in terms of innovation, and, in turn, their investments of R&D and effort generate more spillovers for other firms.

5.3.2. Age Patterns of the Optimal Allocations

As explained in Section 4, age patterns can in principle arise for two reasons: the fact that screening policies are set at age one with full commitment and a possible age-dependency of the primitives of the model. Panels A and B from Figure 2 plot the optimal wedges, averaged over firm type at a given age.

Younger firms simultaneously have their profits taxed at a higher rate (i.e., subsidized at a lower rate here) and their R&D investment expenses on net subsidized less. When types are less than fully persistent (the estimated persistence parameter is 0.63), the screening terms in Proposition 1 are largest in absolute values early in life, when the firm has the most private information, and decay with time, at a rate that is decreasing in the persistence. Hence, it is optimal to distort the allocations more among young firms in order to reduce overall informational rents. Over time, as the screening term decays, the wedges for firms of different productivities converge to the Pigouvian correction and the monopoly valuation term. Because the screening term on the R&D wedge is negative, this means that the net R&D wedge converges to these corrective terms from below, while the profit wedge converges from above.

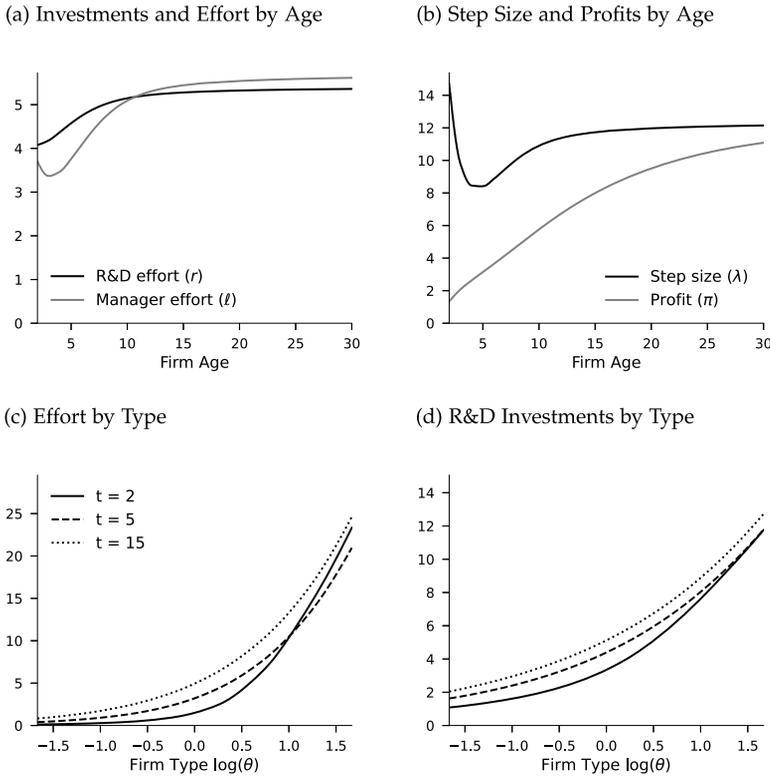


FIGURE 3.—Optimal allocations.

5.4. Comparative Statics: The Role of Persistence, Complementarity, and the Strength of the Spillover

In Figure 4, we quantify the effect of the key parameters. Panels A and B depict the optimal wedges when $\rho_{\theta r} = 0.8 < \rho_{lr} = 1$. When $\rho_{\theta r}$ is smaller, the optimal R&D wedge is larger, especially for lower productivity firms. In this case, since it is not just high productivity firms that benefit from R&D investments, there is no need to reduce the innovation incentives provided to low productivity firms by as much to prevent high productivity firms from mimicking them. In addition, the closer $\rho_{\theta r}$ is to 1, the flatter the net wedge is for different firm ages.

The persistence of the firm’s research productivity process affects the optimal policies very significantly, in particular the age pattern. Panels C and D depict the wedges for a higher value of persistence than our benchmark case, namely, for $\tilde{p} = 0.9$. With a higher persistence, wedges decay at a lower rate. We provide several more comparative statics and robustness checks on the value and shape of the persistence in the replication files, such as a first-order autoregressive process instead of the benchmark logarithmic autoregressive process, an increasing persistence over the life cycle that has the same average value as our benchmark estimate, a lower persistence ($p = 0.5$), and a higher persistence ($p = 0.9$). The speed of convergence to the Pigouvian and monopoly correction terms is strongly shaped by the type process. Yet, although the persistence of this stochastic process affects the rate of decay of the wedges very significantly, it does not change our qualitative findings. In addition, a more persistent process increases the planner’s abil-

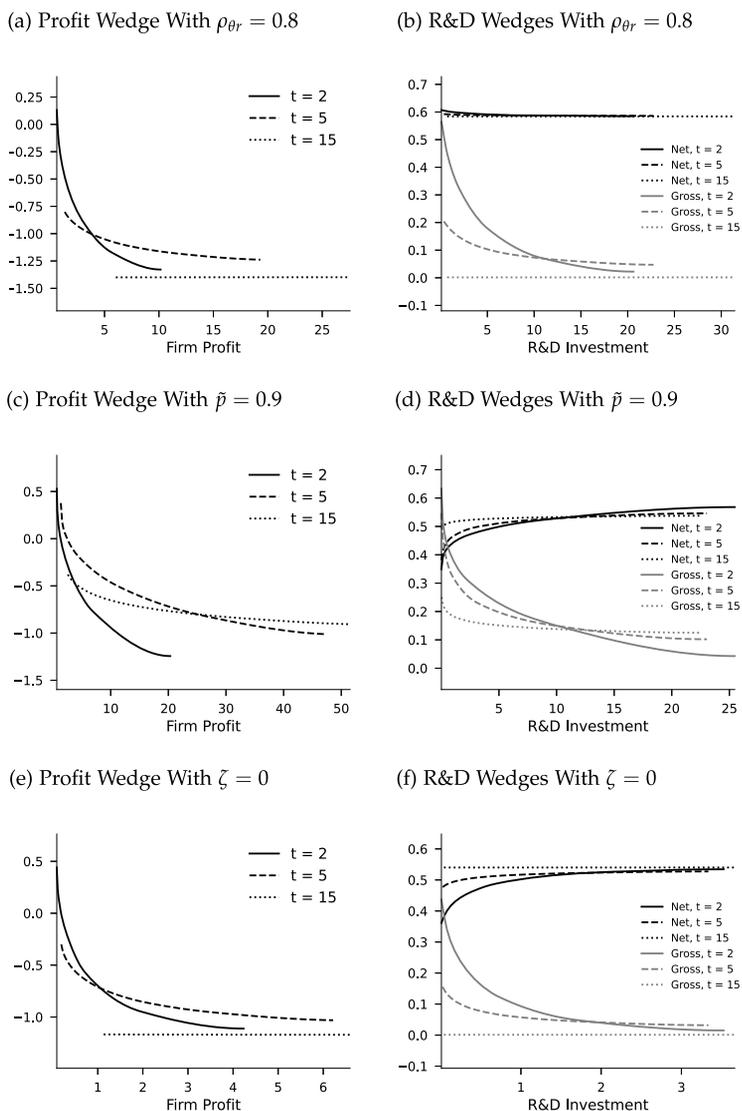


FIGURE 4.—Comparative statics: optimal profit and R&D wedges.

ity to provide dynamic incentives and improves the allocations: there are higher levels of effort and R&D investment for firms of all productivities.

Finally, panels E and F depict the wedges when there is no spillover ($\zeta = 0.0$). In this case, the wedges simply correct for the monopoly distortion. Unsurprisingly, a lack of spillover leads to lower R&D effort and smaller investment wedges.

5.5. Robustness Checks and Sensitivity Analysis

We provide many robustness checks and sensitivity analyses in Section S.5 of the replication file: we perform a type of two-step GMM estimation with weights taken from the variance-covariance matrix of data moments. Furthermore, we explore the role of the stochastic type process assumed, as discussed above. We also vary the value of β , where

higher β represents more market power. At the same time, a higher β also means that the quality of each differentiated product is more valued by consumers. On balance, there is more investment in R&D and more effort at the optimum when β is higher. In addition, we consider higher rates of depreciation of innovation, $\delta = 0.15$ and $\delta = 0.3$. The higher the rate of depreciation, the larger the wedges have to be to induce firms to invest (relative to what they would do if left to choose). Naturally, the higher the rate at which knowledge depreciates, the lower the optimal investments, step sizes, and resulting innovation that can be stimulated. We also show what happens when the cost of R&D is less convex, that is, when $\eta = 1$. This barely changes the wedges, as they represent the share of costs that is subsidized. However, the level of R&D effort and incentives that can be incentivized are larger when costs are less convex. Finally, we re-estimate the model only on publicly-traded firms from COMPUSTAT matched to patent data and show what happens with a finite firm life cycle, in which case the horizon becomes important.

6. SIMPLER INNOVATION POLICIES

Until now, we have considered a fully unrestricted mechanism that does not place constraints other than incentive compatibility on the policies. In this section, we consider restricted, simpler policies. We solve for the optimal policy within each of the restricted classes of policies considered, using the estimated parameters from Section 5. We then compute the welfare loss relative to the welfare obtained with the unrestricted mechanism. Table V shows our results. Each panel considers a separate class of policies, ranging from linear to nonlinear and non-separable policies. The two columns show, respectively, the welfare achieved from the optimal policy in each class relative to (i) our benchmark

TABLE V
WELFARE FROM OPTIMAL SIMPLER POLICIES.

Policy Type	Welfare Achieved Relative to Full Optimum		
	Benchmark	No spillovers	
<i>A. Current US policy</i> $T'(\pi) = 0.23$	$S'(M) = 0.19$	18%	31.1%
<i>B. Optimal Linear</i> $T'(\pi) = \tau_0$	$S'(M) = s_0$	89%	88.5%
<i>C. Linear With Interaction Term</i> $T'(\pi, M) = \tau_0 + \tau_1 M$	$S'(M) = s_0$	93.5%	93.7%
<i>D. Heathcote–Storesletten–Violante (HSV)</i> $T'(\pi) = \tau_0 - \tau_1 \pi^{\tau_2}$	$S'(M) = s_0 - s_1 M^{\tau_2}$	97.4%	98.2%
<i>E. HSV Tax on Profits and Linear Subsidy</i> $T'(\pi) = \tau_0$	$S'(M) = s_0 - s_1 M^{\tau_2}$	94.7%	95.6%
<i>F. HSV Subsidy on R&D and Linear Profit Tax</i> $T'(\pi) = \tau_0$	$S'(M) = s_0 - s_1 M^{\tau_2}$	97.3%	97.4%
<i>G. HSV With Interaction Term</i> $T'(\pi, M) = \tau_0 + \tau_3 M^{\tau_3} - \tau_1 \pi^{\tau_2}$	$S'(M) = s_0 - s_1 M^{\tau_2}$	97.4%	98.3 %

Note: The table shows the share of welfare from the full unrestricted optimum that is achieved by the optimal policy within each class. Each panel shows a different class. Column (1) shows the welfare relative to the benchmark optimum; Column (3) for the benchmark optimum but when there is no spillover ($\zeta = 0$).

optimal planner solution and (ii) the limit case in which there is no spillover ($\zeta = 0$) so that only the monopoly distortion has to be corrected.

The first row shows the welfare level from the current policies in the United States, that is, approximated with a linear 23% effective corporate tax rate and a 19% effective R&D subsidy rate. Current policies only achieve 18% of the gain relative to our benchmark planner problem. If there were no spillovers at all, current policies would do less poorly and achieve 31.1% of the welfare of the optimum.

The next rows show progressively more complex policies. The optimal policy within the linear class with a linear profit tax and a linear subsidy (Panel B) does much better than the current policy and yields 89% of the welfare gain from the full optimum, and 88.5% if there is no spillover. Adding an interaction term between the marginal tax rate and the level of R&D spending (Panel C) further improves welfare gains.

The biggest gain comes from a nonlinear Heathcote–Storesletten–Violante (HSV) policy, as used by Heathcote, Storesletten, and Violante (2017) and Heathcote, Storesletten, and Violante (2020). The HSV policy is a parsimonious parameterized tax function, with one parameter controlling the average level of taxes, and another controlling the progressivity. We extend the HSV policy to allow for a constant component of the marginal tax rate τ_0 , and we parameterize both the profit tax and the R&D subsidy with this HSV-type function. The optimal HSV policy reaps a full 97.4% of the full welfare gain of the optimal mechanism (and up to 98.2% in the case with no spillover). Once this nonlinearity is allowed for, additional nonseparability between profits and R&D expenses brings no further gain (Panel G). The marginal profit tax and the marginal R&D subsidy implied by this HSV function exactly mimic the patterns of the profit and the gross R&D wedges in Panels C and D of Figure 2, with lower marginal taxes (higher marginal profit subsidies) on higher profit firms and lower marginal R&D subsidies for higher levels of R&D investments.

One may further ask whether it is the nonlinearity in the profit tax or the nonlinearity in the subsidy that matters most. We answer this question by simplifying either the subsidy to be linear (Panel E) or the tax to be linear (Panel F) while leaving the other function to be HSV-type as in Panel D. The most important gain comes from a nonlinear R&D subsidy: A linear subsidy plus HSV tax system yields 94.7% of the full gain. On the contrary, linearizing the profit tax generates only a very small welfare loss relative to the fully nonlinear HSV policy and achieves 97.3% of the welfare from the full optimum. Thus, in our estimated model, the most important quantitative features are: first, the nonlinearity in the R&D subsidy that takes a HSV form and which provides lower marginal subsidies for higher levels of R&D investment; second, although the constant profit tax in this case provides excessively generous incentives to the low profitability firms (and just about the right level for high profitability firms), that loss is quantitatively very small since less profitable firms make low profits to start with. This is a particularly useful finding because corporate taxes are typically more or less linear. On the other hand, one can easily imagine a more nuanced HSV-type R&D subsidy scheme being implemented, where the marginal subsidy depends on the investment level.

Does this mean that it is in general optimal to subsidize profits at the margin or have (weakly) lower marginal profit taxes on more profitable firms? Of course, this is the right corporate tax system for innovating firms, and not all firms in the economy are innovating. The reasons to tax or subsidize non-innovating firms would be different. This case could be nested in our model if spillovers are shut off and the model is appropriately calibrated. Our framework for firm taxation is malleable and quite general, and we think it can be used to study firm taxation more broadly. If the government can set different corporate

tax systems based on whether a firm is in an innovating sector, the tax system presented in this paper would apply to the innovative sectors. If the government cannot distinguish between innovating and non-innovating sectors, then the optimal tax system would be a mix of the optimal tax systems for the non-innovating and innovating sectors, allowing for possible shifting between the two. This would be a great avenue for future research, leveraging the methods in this paper to address more complex issues in corporate taxation.

7. CONCLUSION

In this paper, we study how to most efficiently use tax policy to stimulate R&D investments when there are spillovers between firms. Our core contribution is to introduce asymmetric information in a dynamic firm taxation model with spillovers.

With our core setup and methodology in place, additional aspects of R&D investments and innovation by firms can be incorporated, and we discussed some possible generalizations and extensions. Even though we motivate our analysis specifically with R&D investments, our results and the theoretical and numerical solution methods are much more broadly applicable to the provision of firm incentives in dynamic settings with asymmetric information and with other types of investments with or without spillovers. To this end, we wrote our formulas in the most generic form possible. R&D investments are just one of the potential applications of this framework. Our framework provides a new way to think about innovation, but also about firm taxation more generally. Introducing asymmetric information and heterogeneous, stochastic firm types captures many features of the real world and could allow researchers to fruitfully address important questions in policy design for firms.

We hope that future research will build on this fruitful combination of macro-level policy questions, with newly developed mechanism design techniques, which are guided by firm-level micro data, to study many important issues. First, the competition structure in the intermediate goods market could be made endogenous to tax policy: firms would then enter, exit, and steal products from their competitors in response to the tax incentives. Second, one could study optimal R&D policies when there is a noisy signal about product quality that firms may be able to manipulate. Third, a more extended structural estimation focusing on the identification of the key parameters we emphasized (complementarities, persistence, and strength of spillovers) for different sectors and types of products could shed further light on optimal sector-specific policies.

REFERENCES

- ABOODY, D., AND B. LEV (2000): "Information Asymmetry, R&D, and Insider Gains," *The Journal of Finance*, 55 (6), 2747–2766. [656]
- ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): "Innovation, Reallocation and Growth," *American Economic Review*, 108 (11), 3450–3491. [669,673]
- AGHION, P., AND P. HOWITT (1992): "A Model of Growth Through Creative Destruction," *Econometrica*, 60 (2), 323–351. [652]
- AKCIGIT, U., AND W. R. KERR (2018): "Growth Through Heterogeneous Innovations," *Journal of Political Economy*, 126 (4), 1374–1443. [652,658,669,672]
- AKCIGIT, U., D. HANLEY S. STANTCHEVA (2022): "Supplement to 'Optimal Taxation and R&D Policies,'" *Econometrica Supplemental Material*, 90, <https://doi.org/10.3982/ECTA15445>. [648,649]
- ALAM, P., AND K. S. WALTON (1995): "Information Asymmetry and Valuation Effects of Debt Financing," *The Financial Review*, 30 (2), 289–311. [656]
- ALBANESI, S., AND C. SLEET (2006): "Dynamic Optimal Taxation With Private Information," *Review of Economic Studies*, 73 (1), 1–30. [648,668]

- BARON, D. P., AND R. B. MYERSON (1982): "Regulating a Monopolist With Unknown Costs," *Econometrica*, 50 (4), 911–930. [648]
- BLOOM, N., AND R. GRIFFITH (2001): "The Internationalisation of UK R&D," *Fiscal Studies*, 22 (3), 337–355. [649]
- BLOOM, N., AND J. VAN REENEN (2007): "Measuring and Explaining Management Practices Across Firms and Countries," *Quarterly Journal of Economics*, 122 (4), 1351–1408. [649,650]
- BLOOM, N., L. CHENNELLS, R. GRIFFITH, AND J. VAN REENEN (2002): *How Has Tax Affected the Changing Cost of R&D? Evidence From Eight Countries*. Springer. [649]
- BLOOM, N., B. EIFERT, A. MAHAJAN, D. MCKENZIE, AND J. ROBERTS (2013): "Does Management Matter? Evidence From India," *Quarterly Journal of Economics*, 128 (1), 1–51. [649,650]
- BLOOM, N., R. GRIFFITH, AND J. 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. [649,673]
- BLOOM, N., R. SADUN, AND J. VAN REENEN (2012): "The Organization of Firms Across Countries," *Quarterly Journal of Economics*, 127 (4), 1663–1705. [649,650]
- BLOOM, N., M. SCHANKERMAN, AND J. VAN REENEN (2013): "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81 (4), 1347–1393. [649,670,673]
- CHEN, Z., Z. LIU, J. SUÁREZ SERRATO, AND D. XU (2021): "Notching R&D Investment With Corporate Income Tax Cuts in China," *American Economic Review*, 111 (7), 2065–2100. [656]
- FARHI, E., AND I. WERNING (2013): "Insurance and Taxation Over the Life Cycle," *Review of Economic Studies*, 80 (2), 596–635. [648,659,661,668]
- GHARBI, S., J.-M. SAHUT, AND F. TEULON (2014): "R&D Investments and High-Tech Firms' Stock Return Volatility," *Technological Forecasting and Social Change*, 88, 306–312. [656]
- GOLOSOV, M., A. TSYVINSKI, AND I. WERNING (2006): "New Dynamic Public Finance: A User's Guide," in *NBER Macroeconomics Annual*, Vol. 21. MIT Press, 317–388. 2006. [648]
- GOLOSOV, M., A. TSYVINSKI, AND N. WERQUIN (2014): "A Variational Approach to the Analysis of Tax Systems," Working Paper 20780, NBER. [648]
- GOMPERS, P. A. (1995): "Optimal Investment, Monitoring, and the Staging of Venture Capital," *The Journal of Finance*, 50 (5), 1461–1489. [656]
- GOOLSBEE, A. (1998): "Does R&D Policy Primarily Benefit Scientists and Engineers?" *American Economic Review (Papers and Proceedings)*, 88 (2), 298–302. [649]
- GRILICHES, Z. (1998): *R&D and Productivity: The Econometric Evidence*. University of Chicago Press. [673]
- GROSSMAN, V., T. STEGER, AND T. TRIMBORN (2013): "Dynamically Optimal R&D Subsidization," *Journal of Economic Dynamics & Control*, 37, 516–534. [648]
- GUNER, N., G. VENTURA, AND Y. XU (2008): "Macroeconomic Implications of Size-Dependent Policies," *Review of Economic Dynamics*, 11, 721–744. [672]
- HALL, B., AND J. VAN REENEN (2000): "How Effective Are Fiscal Incentives for R&D? A Review of the Evidence," *Research Policy*, 29 (4), 449–469. [656]
- HALL, B., A. JAFFE, AND M. TRAJTENBERG (2001): "The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools," Technical report. [656,669]
- (2005): "Market Value and Patent Citations: A First Look," *RAND Journal of Economics*, 36, 1638. [672]
- HALL, B., AND J. LERNER (2010): "The Financing of R&D and Innovation," in *Handbook of the Economics of Innovation*, Vol. 1, ed. by B. Hall, and N. Rosenberg. Amsterdam: North-Holland, 609–639. [656]
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 132 (4), 1693–1754. [681]
- (2020): "Optimal Progressivity With Age-Dependent Taxation," *Journal of Public Economics*, 189, 104074. [681]
- HICKS, J. (1970): "Elasticity of Substitution Again: Substitutes and Complements," *Oxford Economic Papers*, 25, 289–296. [651]
- JONES, C. I., AND J. C. WILLIAMS (1998): "Measuring the Social Return to R&D," *Quarterly Journal of Economics*, 113 (4), 1119–1135. [649]
- (2000): "Too Much of a Good Thing? The Economics of Investment in R&D," *Journal of Economic Growth*, 5, 65–85. [649]
- KAPLAN, S. N., AND P. STROMBERG (2001): "Venture Capitalists as Principals: Contracting, Screening, and Monitoring," *The American Economic Review*, 91 (2), 426–430. [656]
- KERR, W. R., R. NANDA, AND M. RHODES-KROPP (2014): "Entrepreneurship as Experimentation," *Journal of Economic Perspectives*, 28 (3), 25–48. [656]
- LAFFONT, J.-J., AND J. TIROLE (1986): "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 94 (3), 614–641. [648,654]

- LOCKWOOD, B. B., C. G. NATHANSON, AND E. GLEN WEYL (2017): "Taxation and the Allocation of Talent," *Journal of Political Economy*, 125 (5), 1635–1682. [648]
- MANSFIELD, E. (1986): "The R&D Tax Credit and Other Technology Policy Issues," *The American Economic Review*, 76 (2), 190–194. [656]
- MILGROM, P., AND I. SEGAL (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70 (2), 583–601. [659]
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 82 (2), 601–653. [648,659]
- PIKETTY, T., E. SAEZ, AND S. STANTCHEVA (2014): "Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities," *American Economic Journal: Economic Policy*, 6 (1), 230–271. [648]
- ROMER, P. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98 (5), 71–102. [652]
- ROTHSCHILD, C., AND F. SCHEUER (2013): "Redistributive Taxation in the Roy Model," *The Quarterly Journal of Economics*, 128 (2), 623–668. [648]
- (2016): "Optimal Taxation With Rent-Seeking," *The Review of Economic Studies*, 83 (3), 1225–1262. [648]
- SACHS, D., A. TSYVINSKI, AND N. WERQUIN (2020): "Nonlinear Tax Incidence and Optimal Taxation in General Equilibrium," *Econometrica*, 88, 469–493. [648]
- SAPPINGTON, D. (1982): "Optimal Regulation of Research and Development Under Imperfect Information," *Bell Journal of Economics*, 13 (2), 354–368. [648]
- STANTCHEVA, S. (2015): "Learning and (or) Doing: Human Capital Investments and Optimal Taxation," Working Paper 21381, National Bureau of Economic Research. [648]
- (2017): "Optimal Taxation and Human Capital Policies Over the Life Cycle," *Journal of Political Economy*, 125 (6), 1931–1990. [648,659,661,668]
- TYSON, L., AND G. LINDEN (2012): *The Corporate R&D Tax Credit and U.S. Innovation and Competitiveness: Gauging the Economic and Fiscal Effectiveness of the Credit*. Washington: Center for American Progress. [650]
- WERQUIN, N. (2016): "Income Taxation With Frictional Labor Supply," Working Paper, Toulouse School of Economics. [648]
- ZANTOUT, Z. Z. (1997): "A Test of the Debt-Monitoring Hypothesis: The Case of Corporate R&D Expenditures," *The Financial Review*, 32 (1), 21–48. [656]

Co-editor Giovanni L. Violante handled this manuscript.

Manuscript received 19 June, 2017; final version accepted 19 August, 2021; available online 27 September, 2021.

SUPPLEMENT TO “OPTIMAL TAXATION AND R&D POLICIES”
 (Econometrica, Vol. 90, No. 2, March 2022, 645–684)

UFUK AKCIGIT
 Economics, University of Chicago, NBER, and CEPR

DOUGLAS HANLEY
 Economics, University of Pittsburgh

STEFANIE STANTCHEVA
 Economics, Harvard, NBER, and CEPR

APPENDIX OA.1: PROOFS OF THE PROPOSITIONS IN THE MAIN TEXT

PROOF OF PROPOSITION 1: TAKING THE FOC of program P in (8) with respect to $r_t(\theta^t)$ yields

$$\begin{aligned}
 [r_t(\theta^t)]: \frac{1}{R} \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \frac{\partial \tilde{Y}(\theta^s, \bar{q}_s)}{\partial q_s} \frac{\partial \lambda(\theta^{t+1})}{\partial r_t} \right) \\
 - \frac{1}{R} \mathbb{E} \left(\frac{1-F^1(\theta_1)}{f^1(\theta_1)} I_{1,t+1}(\theta^{t+1}) \phi'_{t+1}(l(\theta^{t+1})) \frac{\lambda_{\theta} \lambda_r}{\lambda \lambda_l} [\rho_{\theta r} - \rho_{lr}] \right) \\
 - M'_t(r(\theta^t)) + \frac{1}{R} \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \eta_s \frac{\partial \lambda(\theta^{t+1})}{\partial r_t} \right) = 0.
 \end{aligned}$$

Using the definition of the R&D wedge as

$$s(\theta^t) = M'_t(r(\theta^t)) - \frac{1}{R} \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \frac{\partial \pi_s(\theta^s)}{\partial q_s} \frac{\partial \lambda_{t+1}}{\partial r_t} \right)$$

to substitute for the marginal cost $M'_t(r_t(\theta^t))$ in the FOC, we obtain formula (10).

Taking the FOC with respect to $l_t(\theta^t)$ yields

$$\begin{aligned}
 [l_t(\theta^t)]: \mathbb{E} \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \tilde{Y}(\theta^s, \bar{q}_s)}{\partial q_s} \frac{\partial \lambda(\theta^t)}{\partial l_t} \right) \\
 - \frac{1-F^1(\theta_1)}{f^1(\theta_1)} I_{1,t}(\theta^t) \frac{\partial}{\partial l_t} \left[\phi'_t(l_t(\theta^t)) \frac{\partial \lambda(\theta^t) / \partial \theta_t}{\partial \lambda(\theta^t) / \partial l_t} \right] \\
 - \phi'_t(l_t(\theta^t)) + \mathbb{E} \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \eta_s \frac{\partial \lambda(\theta^t)}{\partial l_t} \right) = 0.
 \end{aligned}$$

Ufuk Akcigit: uakcigit@uchicago.edu
 Douglas Hanley: doughanley@pitt.edu
 Stefanie Stantcheva: sstantcheva@fas.harvard.edu

Transform the derivative of the envelope condition (using the notation ϕ_{ll} and $\phi_{ll,t}$ to denote, respectively, the first and second derivatives of ϕ with respect to l):

$$\begin{aligned} \frac{\partial}{\partial l_t} \left[\phi_{ll} \frac{\lambda_{\theta t}}{\lambda_{ll,t}} \right] &= \left(\phi_{ll,t} - \phi_{ll} \frac{\lambda_{ll,t}}{\lambda_{ll,t}} \right) \frac{\lambda_{\theta t}}{\lambda_{ll,t}} + \phi_{ll} \frac{\lambda_{\theta l,t}}{\lambda_{ll,t}} = \frac{\phi_{ll} \lambda_{\theta t}}{\lambda_{ll,t}} \left[\frac{\left(\phi_{ll,t} - \phi_{ll} \frac{\lambda_{ll,t}}{\lambda_{ll,t}} \right) \lambda_{ll,t}}{\phi_{ll}} \frac{\lambda_{\theta t}}{\lambda_{ll,t}} + \frac{\lambda_{\theta l,t} \lambda_{ll,t}}{\lambda_{\theta t} \lambda_{ll,t}} \right] \\ &= \frac{\phi_{ll} \lambda_{\theta t}}{\lambda_{ll,t}} \left[\frac{1}{\varepsilon_{l,1-\tau}} \frac{\lambda_{ll,t}}{\lambda_{ll,t}} + \rho_{\theta l,t} \right] = \frac{\phi_{ll} \lambda_{\theta t}}{\lambda_{ll,t}} \left[\frac{1}{\varepsilon_{l,1-\tau}} \frac{1}{\varepsilon_{\lambda l,t}} + \rho_{\theta l,t} \right]. \end{aligned}$$

Using the definition of the wedge $\tau(\theta^t)$ to substitute for $\phi'_l(l_t(\theta^t))$ yields the formula in the text.

PROOF OF IMPLEMENTATION RESULT: For every period, define the following objects:

$$\begin{aligned} D_s(\theta^{s-1}, \theta_s) &= E \left(\sum_{t=s}^{\infty} I_{(s),t} \left(\frac{1}{R} \right)^{t-s} \frac{\partial v_t}{\partial \theta_t} \middle| \theta^s \right), \\ Q_s(\theta^{s-1}, \theta_s) &= \int_{\underline{\theta}}^{\theta_s} D_s(\theta^{s-1}, q) dq, \end{aligned}$$

where the expectation is explicitly conditioned on history θ^t .

With a stochastic process such that the impulse response is independent of θ^t except through θ_1 and θ_t , we have that $I_{(s),t} = i(\theta_1, \theta_t, t)$ for some function $i(\cdot)$. In addition, $\frac{\partial v_t}{\partial \theta_t} = \phi'_t(l_t(\theta^t)) \frac{\partial \lambda(\theta^t)}{\partial \theta_t}$, so that

$$D_s(\theta^{s-1}, \theta_s) = E \left(\sum_{t=s}^{\infty} \left(\frac{1}{R} \right)^{t-s} i(\theta_1, \theta_t, t) \phi'_t(l_t(\theta^t)) \frac{\partial \lambda(\theta^t) / \partial \theta_t}{\partial \lambda(\theta^t) / \partial l_t} \middle| \theta^s \right).$$

In the unrestricted mechanism, the transfers provided every period are

$$T_t(\theta^t) = Q_t(\theta^{t-1}, \theta_t) - \frac{1}{R} E_t(Q_{t+1}(\theta^t, \theta_{t+1})) + \phi_t(l_t(\theta^t)). \quad (\text{OA1})$$

Given the time separable utility and the assumption on the impulse response functions, the transfer hence depends on λ_t , r_{t-1} , θ_t , and θ_1 (and, naturally, on age t). Denote it by $T_t^*(\lambda_t, r_{t-1}, \theta_t, \theta_1)$.

With the price subsidy in place, the total price faced by the monopolist is $\frac{Y(q,k)}{k}$. Hence, conditional on q_t , the monopolist maximizes social surplus from production and the choice will be a deterministic function of quality, denoted by $k_t(q_t)$. As a result, profits earned are a deterministic function of quality, denoted by $\pi_t(q_t)$.

Note that in period 1, since r_0 and q_0 are given and observed, the realization

$$q_1 = H(q_0, \lambda_1(l(\theta_1), r(\theta_0), \theta_1))$$

can be inverted to obtain θ_1 (at the optimal allocation, under incentive compatibility) as long as for every θ_1 there is a uniquely optimal $l(\theta_1)$. Hence, we will use conditioning on q_1 instead of θ_1 . Let $\Theta^t(q_1, r_{t-1}, q_{t-1})$ be the set of all histories (including θ_t) that

are consistent with q_1 in period 1, and r_{t-1} and q_{t-1} . For each θ_t in this set, the optimal allocations and transfer are the same (independent of what exactly happened in the full past). Let $r_t^*(\theta)$, $l_t^*(\theta)$ be the optimal allocations given to each θ in this set (they are equal for each such θ by inspection of the wedge formulas at the optimum). The implied optimal quality is then $q_t^*(\theta) = q_{t-1} + \lambda_t(r_{t-1}, l_t^*(\theta), \theta)$.

We now have to make the tax system such that allocations which do not arise in the Planner's solution are very unattractive to the agent. First, we can rule out allocations that never occur for any θ in $\Theta^t(q_1, r_{t-1}, q_{t-1})$ by making the transfer at points $q_t^*(\theta)$, $r_t^*(\theta)$ following q_{t-1} , r_{t-1} , q_1 highly negative. We can also directly rule out histories q_{t-1} and r_{t-1} which should never occur in the Planner's problem in the same way.

For all remaining consistent histories and for each θ in $\Theta^t(q_1, r_{t-1}, q_{t-1})$, the tax or transfer given as a function of the observables needs to be such that

$$T_t(q_t^*(\theta), r_t^*(\theta), q_{t-1}, r_{t-1}, q_1) + \pi_t(q_t^*(\theta)) = T_t^*(\lambda_t(r_{t-1}, l_t^*(\theta), \theta), r_{t-1}, \theta).$$

Consider the firm's choice. First, for given r_{t-1} , q_{t-1} , and θ_t , the firm should rationally only select a pair q_t^* , r_t^* that is consistent with some $\theta \in \Theta^t(q_1, r_{t-1}, q_{t-1})$ or else the transfer it receives would be very negative. For each r_{t-1} , q_{t-1} , and θ_t , if the firm chooses $q_t^*(\theta)$ and $r_t^*(\theta)$ meant for type θ in the planner's problem, it receives the utility it would get from reporting to be type θ in the Planner's problem. By incentive compatibility, the firm will choose the allocation meant for its true type realization. *Q.E.D.*

APPENDIX OA.2: WORKED EXAMPLE WITH CONSTANT MARKUPS

Production

We can specialize the functional form to one that delivers constant markups. Let the cost of production be $C(k, \bar{q}) = \frac{k}{\bar{q}^\xi}$, and the output as valued by consumers be $Y(q_t(\theta^t), k_t(\theta^t)) = \frac{1}{1-\beta} q_t(\theta^t)^\beta k_t(\theta^t)^{1-\beta}$. The demand function under a patent system that grants monopoly rights is then

$$p(q_t(\theta^t), k_t(\theta^t)) = q_t(\theta^t)^\beta k_t(\theta^t)^{-\beta}$$

and the quantity chosen by the monopolist is

$$k(q_t(\theta^t), \bar{q}_t) = [(1-\beta)\bar{q}_t^\xi]^{-\frac{1}{\beta}} q_t(\theta^t).$$

At the optimum, the price is a constant markup over marginal cost equal to

$$p(\bar{q}_t) = \frac{1}{(1-\beta)\bar{q}_t^\xi}.$$

Profits are then given by

$$\pi(q_t(\theta^t), \bar{q}_t) = q_t(\theta^t)(1-\beta)^{\frac{1-\beta}{\beta}} \cdot \beta \cdot \bar{q}_t^\xi \frac{1-\beta}{\beta}.$$

$Y(q_t(\theta^t), \bar{q}_t)$, the output from the private producer in the laissez-faire with a monopoly right, is

$$Y(q_t(\theta^t), \bar{q}_t) = Y(q_t(\theta^t), k(q_t(\theta^t), \bar{q}_t)) = \frac{1}{1-\beta} q_t(\theta^t) ((1-\beta)\bar{q}_t^\xi)^{\frac{1-\beta}{\beta}}.$$

Hence, the final good in the private market equilibrium is given by

$$Y_t = \int_{\theta^t} Y(q_t(\theta^t), \bar{q}_t) P(\theta^t) = \int_{\theta^t} \frac{1}{1-\beta} q_t(\theta^t) [(1-\beta)\bar{q}_t^\zeta]^{\frac{1-\beta}{\beta}} P(\theta^t) d\theta^t.$$

Conditional on a given quality $q_t(\theta^t)$, the production choice of the planner would be such that

$$k^*(q_t(\theta^t), \bar{q}_t) = \bar{q}_t^{\frac{\zeta}{\beta}} q_t(\theta^t) > k(q_t(\theta^t), \bar{q}_t).$$

A Special Case With Very Simple Wedges

We can impose additional restrictions to obtain particularly easy characterizations of the wedges. Assume the functional forms in Table II, but also assume the special case in which $\rho_{\theta^r} = \rho_{r^t} = 1$, so that the screening term in the R&D wedge is zero.

Let

$$B_e = 1 + \zeta \left(\frac{1-\beta}{\beta} \right),$$

$$B_m = \frac{2-\beta}{1-\beta},$$

and

$$G_t = H^1(\theta^1) p^{t-1} (1+\gamma)(1-\alpha).$$

Then, we can show that in this special case,

$$\frac{\tau_t}{1+G_t} = -\left(1 - \frac{1}{B_e}\right) - \frac{1}{B_e} \left(1 - \frac{1}{B_m}\right),$$

$$s_t = \left(1 - \frac{1}{B_e}\right) + \frac{1}{B_e} \left(1 - \frac{1}{B_m}\right),$$

and so the profit wedge τ_t depends only on time t and the initial state θ_1 and tends to a constant profit subsidy $-(1 - \frac{1}{B_e}) - \frac{1}{B_e}(1 - \frac{1}{B_m}) < 0$ over time. The net subsidy wedge is constant over time and type and equal to exactly $-\tau_t$. Both wedges are increasing in absolute value when the strength of the spillover (ζ) increases.

APPENDIX OA.3: EXTENSIONS

OA.1. *Heterogeneity in Production Efficiency*

Suppose that firms are also heterogeneous in their production productivities, denoted by θ^p , with realization θ_i^p and history $\theta^{p,t}$. For instance, production costs could be $C(k, \bar{q}_t, \theta_i^p)$. Allocations are now specified as functions of the full set of histories $(\theta^t, \theta^{p,t})$. If production productivity is observable, the planner will simply condition on it for each history of research productivities θ^t . In fact, as long as quality q and quantity k are observable, the planner can perfectly infer $\theta^{p,t}$ from the observed production choices. Net output is then $\tilde{Y}(q_t(\theta^t, \theta^{p,t}), \bar{q}_t, \theta_i^p)$ and profits are $\pi(q_t(\theta^t, \theta^{p,t}), \bar{q}_t, \theta_i^p)$. Similarly to before, we can define $\Pi_t(\theta^t, \theta^{p,t}) := (\sum_{s=t}^{\infty} (\frac{1-\delta}{R})^{s-t} \frac{\partial \pi(q_s(\theta^s, \theta^{p,s}), \bar{q}_s, \theta_i^p)}{\partial q_s})$ as the marginal impact on future

profit flows from an increase in quality. Let $Q_t(\theta^t, \theta^{p,t}) = \sum_{s=t}^{\infty} \left(\frac{1-\delta}{R}\right)^{s-t} \frac{\partial \bar{Y}(q_s(\theta^s, \theta^{p,s}), \bar{q}_s, \theta_s^p)}{\partial q_s}$ be the marginal impact of quality on future expected output net of production costs.

Then, the optimal profit wedge can be set for each history $(\theta^t, \theta^{p,t})$ and satisfies

$$\begin{aligned} \tau(\theta^t, \theta^{p,t}) = & -\mathbb{E} \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \eta_s \right) \frac{\partial \lambda_t}{\partial l_t} - \mathbb{E} (Q_t(\theta^t, \theta^{p,t}) - \Pi_t(\theta^t, \theta^{p,t})) \frac{\partial \lambda_t}{\partial l_t} \\ & + \frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t}(\theta^t) \frac{\phi'_t \lambda_{\theta t}}{\lambda_t} \left[\frac{1}{\varepsilon_{l,1-\tau}} \frac{1}{\varepsilon_{\lambda,t}} + \rho_{\theta l,t} \right] \end{aligned}$$

and the optimal R&D subsidy is given by

$$\begin{aligned} s(\theta^t, \theta^{p,t}) = & \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \eta_s \frac{\partial \lambda_{t+1}}{\partial r_t} \right) \\ & + \mathbb{E} \left((Q_{t+1}(\theta^{t+1}, \theta^{p,t+1}) - \Pi_{t+1}(\theta^{t+1}, \theta^{p,t+1})) \frac{\partial \lambda(\theta^{t+1})}{\partial r_t} \right) \\ & + \frac{1}{R} \mathbb{E} \left(\frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t+1}(\theta^{t+1}) \phi'_{t+1}(l(\theta^{t+1})) \frac{\lambda_{\theta} \lambda_r}{\lambda \lambda_l} (\rho_{lr} - \rho_{\theta r}) \right). \end{aligned}$$

The productivity differences only enter the monopoly valuation term, as they only affect how effectively each firm can transform the quality into output. As a result, productivity differences in production do not really change the previous results.

More generally, any additional heterogeneity that is observable can be treated in a similar way, by conditioning the optimal policies on it. The problem becomes much more complicated if there is additional unobservable heterogeneity that is correlated with research productivity θ . Already in much simpler static settings without spillovers, [Rochet and Choné \(1998\)](#) showed that, with two-dimensional heterogeneity, there are barely any general results. Incorporating non-trivial two-dimensional heterogeneity in a dynamic model with spillovers like this one (and being able to estimate it) would be an important big step for future research.

Empirically, we do not let this additional observable heterogeneity (such as production sector, technology sector, or business-cycle induced effects) contaminate the results and filter it out from the variables thanks to fixed effects before computing our data moments. What could be quite interesting for future research would be to actually specifically estimate the model and simulate differentiated optimal policies, allowing explicitly for different sectors, different technology classes, or different parts of the business cycles.

OA.2. Different Types of Observable R&D Investments

Suppose that there are several types of observable R&D investments that firms can make, denoted by $r^1, \dots, r^j, \dots, r^J$. A natural interpretation would be the investments in different technology classes.

The step size is determined as a function of the observable R&D investments, unobservable R&D effort, and firm research productivity:

$$\lambda_t = \lambda_t(r_{t-1}^1, \dots, r_{t-1}^j, \dots, r_{t-1}^J, l_t, \theta_t).$$

We can define the Hicksian complementarity of each R&D type with firm effort and research productivity as

$$\rho_{\theta r, t}^j := \frac{\frac{\partial^2 \lambda_t}{\partial r_{t-1}^j \partial \theta_t} \lambda_t}{\frac{\partial \lambda_t}{\partial \lambda_t} \frac{\partial \lambda_t}{\partial \theta_t} \frac{\partial \lambda_t}{\partial r_{t-1}^j}} \quad \text{and} \quad \rho_{l r, t}^j := \frac{\frac{\partial^2 \lambda_t}{\partial l_{t-1} \partial l_t} \lambda_t}{\frac{\partial \lambda_t}{\partial l_t} \frac{\partial \lambda_t}{\partial r_{t-1}^j}}.$$

Different types of R&D investments can have very different complementarity profiles with R&D effort and firm type (or, equivalently, their exposure to risk as embodied by the stochastic type). Some investments may generate returns with high certainty, regardless of the type realization, while others may only yield returns when firms are particularly good or in period of good realizations of the stochastic type.

Let the subsidy on investment r_t^j be denoted by $s^j(\theta^t)$. At the optimum, formula (10) holds separately for each type of R&D investment wedge $s^j(\theta^t)$. The wedge $s^j(\theta^t)$ will be increasing in the effect of investment j on the step size (in the Pigouvian correction term), as well as in the relative complementarity of that investment to unobservable R&D effort relative to its complementarity with respect to firm research productivity, $\rho_{ol}^j - \rho_{or}^j$.

The lesson is that, while it is optimal to subsidize investments with higher externalities at a higher rate, it is not as beneficial if these investments are also highly sensitive to the firm productivity and firm research productivity is unobservable.

OA.3. Different Externalities From Different Types of Research

It is also possible to directly incorporate different externalities from each type of R&D investments by letting the cost function be decreasing in each aggregate investment type:

$$C(k, \bar{q}^1, \dots, \bar{q}^J) \quad \text{with} \quad \bar{q}^j = \int_{\Theta^t} q_t^j(\theta^t) d\theta^t \quad \text{and} \quad q_t^j(\theta^t) = q_t^j(\theta^{t-1})(1 - \delta) + \lambda_t^j(r_{t-1}^j, l_t, \theta_t).$$

This is important in order to be able to speak to the very different spillovers from different types of research such as basic and applied research. Basic research may only add little to the total quality of a firm's product, but if its effect on the costs of production of other firms is important, it will suffer from a large under-investment in the laissez-faire, as highlighted in [Akcigit, Hanley, and Serrano-Velarde \(2021\)](#), and will warrant a large Pigouvian correction.

At the firm level, the (single) product quality is given by

$$q_t = (1 - \delta)q_{t-1} + \sum_{j=1}^J \lambda_t^j(r_{t-1}^j, l_t, \theta_t).$$

We have to impose j consistency constraints in the partial program in each period t , each with multiplier η_t^j . Formula (10) then tells us that R&D investments with the highest spillovers (highest $\eta_t^j = \int_{\Theta^t} \frac{\partial \bar{Y}^*(q_t(\theta^t), \bar{q}_t^1, \dots, \bar{q}_t^J)}{\partial \bar{q}_t^j} P(\theta^t) d\theta^t$) are the ones that should be subsidized most (bearing in mind that their complementarities with effort and firm research productivity may dampen the benefits from subsidizing them).

APPENDIX OA.4: COMPUTATIONAL APPENDIX

OA.1. *Computational Procedure*

All code is written in standard Python 3, and depends only on common numerical and scientific modules such as numpy, scipy, pandas, statsmodels, patsy, and matplotlib. The parameter estimation and optimal policy calculations are done using either the Nelder–Mead algorithm or simulated annealing.

Because of the staggered nature of research spending and firm effort decisions, we find the optimal decisions for a log-uniform grid of possible (θ_t, θ_{t+1}) values. In addition, in the case of the optimal mechanism, one also tracks the initial type θ_1 , as this bears on the constraints imposed by informational limitations.

When solving for both the optimal mechanism and the linear tax equilibrium outcome, the solution method is constructed as a fixed point problem on the path of \bar{q} . Because \bar{q} evolves according to a firm’s research decisions and these decisions are made based on expectations that condition on the future path of \bar{q} , the decisions made by firms are in a sense both forward and backward looking.

Given a certain candidate path for \bar{q} , we can find the optimal choices for research spending and firm effort (for either the firm or the planner), which itself amounts to solving a one-dimensional equation for each point in the type space in each time period. Using these decisions, one can construct an updated path for \bar{q} . When this process reaches a fixed point, we have found the equilibrium path for \bar{q} . In practice, as the equations characterizing firm choices are analytical but not closed form, it is more efficient to formulate the problem as a fixed point over both the path of \bar{q} and firm choices for r and l for each type. Updating is then done only using the $M'(r)$ and $\phi'(l)$ terms in the first-order conditions. Additionally, it is useful to dampen the updating process to avoid any numerical instabilities.

Moving to nonlinear policies considerably complicates matters. In this case, the relevant state space of the firm must include the actual value of q . As a result, we must track the joint distribution of q_t , θ_t , and θ_{t+1} . Conceptually, the convergence process and criterion are similar to the linear case, but the run time is much longer. The advantage is that we can entertain tax and subsidy policies that are arbitrary (differentiable) functions of firm profit and R&D investment.

To generate simulated moments for parameter estimation, we simulate a large number of firms ($2^{15} = 32,768$) for the entirety of their life cycle and compute various statistics on this panel of simulated data. All of the moments are relatively straightforward to calculate, with the notable exception of the coefficients for the spillover regression (M8) and the R&D-cost elasticity regression (M9), which are used to identify the externality parameter and various cost elasticities.

For the spillover regression (M8), we actually re-solve and re-simulate the model for a variety of different scenarios in which innovations contribute an additional boost to average productivity \bar{q} , which we interpret as innovation spillovers between firms. We perform this exercise for a variety of boost parameters centered around unity (the baseline model value). We interpret each simulated economy as representing a particular industry with a particular level of innovation spillovers. This mimics the exogenous variation used to identify the spillovers in the [Bloom, Schankerman, and Van Reenen \(2013\)](#) paper. Using this variation, we then run a regression of firm sales on the amount of research spending undertaken by the firm as well as the average research spending by all firms in that time period and industry. We then match this to an analogous regression run by [Bloom, Schankerman, and Van Reenen \(2013\)](#).

Similarly, for the R&D-cost regression (M9), we simulate a variety of economies having different values of the R&D cost parameters κ_r , centered at the baseline value. These differences can represent actual differences in cost, or alternatively, differences in R&D subsidy levels or tax credits. We then run a firm-level regression across time and industry of R&D investment on the level of κ_r .

To generate estimates for the standard errors of our parameter estimates, we take 100 draws from the distribution induced by our data moment means and variances, fully re-estimate the parameters of our model for each of these draws, then report the standard deviation of these estimates. Because some of our data moments (in particular, moments M8 and M9) come from different sources, it is not clear what the interpretation of off-diagonal elements would be. A natural choice is to set them to zero, using a diagonal matrix for the data moment standard errors.

OA.2. *Ex Post Verification Procedure*

To perform the ex post verification, we start with the allocations under truth-telling in the optimal mechanism, $\lambda(\theta^t)$, $r(\theta^t)$, and $T(\theta^t)$ (where the transfers $T(\theta^t)$ are constructed following (OA1)). These allocations are defined for all histories θ^t which could arise along the equilibrium path by the optimal mechanism—thus, any history θ^t that can never arise given the distribution of stochastic shocks is ruled out (with, for instance, infinitely negative transfers $T(\theta^t)$).

For every history θ^{t-1} , we can compute the allocations that would be assigned to an agent of type θ who reports θ' (not necessarily truthfully) among the feasible types in the space Θ at time t . Under any report θ' , the agent will be assigned the allocations $\lambda(\theta^{t-1}, \theta')$, $r(\theta^{t-1}, \theta')$ and $T(\theta^{t-1}, \theta')$, which are meant for the “true” type (θ^{t-1}, θ) . The agent whose true type realization is θ chooses the report θ' that will maximize his expected discounted payoff, which is

$$\max_{\theta'} T(\theta^{t-1}, \theta') - \phi(\lambda(\theta^{t-1}, \theta')/w(r_{t-1}(\theta^{t-1}), \theta)) + \frac{1}{R} \int \omega(\theta^{t-1}, \theta', \theta_{t+1}) f^{t+1}(\theta_{t+1}|\theta).$$

The ex post verification consists in checking whether the agent will, in fact, choose $\theta' = \theta$ (i.e., report his true type) when faced with the set of allocations that can arise for *any* type at the optimum. Note that this amounts to checking that the global incentive constraints are satisfied at the optimal allocations derived using the first-order approach.

REFERENCES

- AKCIGIT, U., D. HANLEY, AND N. SERRANO-VELARDE (2021): “Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth,” *The Review of Economic Studies*, 88 (1), 1–43. [6]
 BLOOM, N., M. SCHANKERMAN, AND J. VAN REENEN (2013): “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 81 (4), 1347–1393. [7]
 ROCHET, J.-C., AND P. CHONÉ (1998): “Ironing, Sweeping, and Multidimensional Screening,” *Econometrica*, 66 (4), 783–826. [5]

Co-editor Giovanni L. Violante handled this manuscript.

Manuscript received 19 June, 2017; final version accepted 19 August, 2021; available online 27 September, 2021.