

Optimal Taxation and R&D Policies*

Ufuk Akcigit Douglas Hanley Stefanie Stantcheva

November 19, 2019

Abstract

We study the optimal design of corporate taxation and R&D policies as a dynamic mechanism design problem with spillovers. Firms are heterogeneous in their research productivity, i.e., in the efficiency with which they convert a given set of R&D inputs into successful innovations and that research productivity is private information. There are non-internalized technological spillovers across firms, but the asymmetric information prevents 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 – that provides lower marginal subsidies at higher R&D levels – can do almost as well as the full 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.

JEL Codes: H21, H23, H25, H32

*Akcigit: University of Chicago, NBER, and CEPR (e-mail: uakcigit@uchicago.edu); Hanley: University of Pittsburgh (e-mail: doughanley@pitt.edu); Stantcheva: Harvard, NBER, and CEPR (e-mail: 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.

1 Introduction

There are many policies the government can potentially use to foster innovation: improve competition, regulate the intellectual property rights regime, directly fund and perform R&D in public institutions, and use taxation and subsidy policies. In this paper we focus on the last of these tools and consider the optimal design of taxation and R&D policies under asymmetric information. We use new methods from the public economics literature, theoretical advances of mechanism design, and firm-level data matched to patent data to discipline and quantify our analysis.

There are two main motivations for our study. First, R&D policies are widespread, not fully understood, and very costly. The U.S. spent 10.8 billion USD on the R&D tax credit in 2012, and 50.56 billion USD on contracting with non Federally Funded Research and Development Centers (FFRDCs) (Tyson and Linden, 2012). Governments all over the world already intervene heavily in the innovation process of private businesses. The share of private business R&D spending that is shouldered by the government is very high in many countries: in the U.S., it is 14%, while in France or Canada, it is close to 25%.¹

Not only do governments intervene in the private R&D and innovation process, they do so through a very wide variety of policies, including, but not limited to, tax credits, tax deductions, direct grants for research, contracting with private firms, subsidies for R&D costs, or direct funding in FFRDCs. The configurations of these many policies also vary widely.² Many countries have size-dependent policies, through which small businesses are treated more favorably, for instance the Small Business Innovation Research (SBIR) program in the U.S. Policies sometimes depend on firm age, e.g., to encourage new firm creation, as through the start-up credit in the U.S. The sheer scale of public resources spent on R&D and the variety of the policies thus funded raises the question of what the right design of R&D policies should be.

Second, asymmetric information is a key feature in the innovation arena. The innovation literature has extensively addressed how to deal with spillovers, but it has much less dealt with asymmetric information about firms and how to screen firms that are good at innovation from bad ones. Yet, as a large empirical literature, reviewed in Section 2.4 shows, the quality of a firm's organization, management, processes or ideas – which shape its innovation outcomes, conditional on inputs – are private information and very difficult for outside parties, including the government, to observe.³ The literature documents the manifestations of asymmetric information in innovation, between firms and their shareholders or investors; this problem is even more pronounced between firms and the government. In addition, in our data, if one tries to

¹Source: OECD R&D Tax Incentive Indicators (available at www.oecd.org/sti/rd-tax-stats.htm) and OECD, National Accounts and Main Science and Technology Indicators, “direct government funding of business R&D and tax incentives for R&D” table.

²For instance, some tax credits, as in the U.S., are computed based on the growth in R&D spending relative to some base level of past R&D (previously a moving average of the company's past investments, now an average over a fixed period), others, as in France, are computed partially on the increment and partially on the absolute level of spending.

³Empirically, the importance of firm heterogeneity and management quality as determinants of firm productivity has been vividly highlighted in a series of key papers (Bloom and Van Reenen (2007), Bloom, Sadun, and Van Reenen (2012), Bloom et al. (2013)).

predict the innovation quality of a firm, the prediction is very poor. Seeing which firms are good at innovation and which ones are not is inherently difficult, even using a very large set of observables that is likely a very generous overestimate of what the government could realistically condition policies on. One way to address the asymmetric information problem is that adopted by Venture Capitalist firms, which perform very hands-on and thorough screening, and provide staged financing subject to intense monitoring. But this intensive hands-on approach is not easily scalable and thus not applicable when it comes to thinking about large-scale government policies.

In this paper, we build a framework that captures this essential aspect of asymmetric information in innovation and addresses the following questions theoretically and quantitatively: Without restricting the set of policy tools a priori, what are the best policies for innovation? What key parameters do optimal policies depend on? Are there simple policies that are almost as good as the fully optimal ones?

There are two market failures in our setting 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. IPR policy may, however, itself create a distortion, as is the case for instance of a patent system that grants firms monopoly rights.

The key feature of our analysis – and the main impediment to fixing the market distortions in a non-distortionary way – is that firms are heterogeneous in their research productivity and, importantly, this research productivity is private information and unobservable to the government. A higher research productivity allows a firm to convert a given set of research inputs into a better innovation output. In addition, while some of the inputs into the R&D process are observable (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 advance information about its future productivity, it cannot perfectly foresee it. As a result, at the time when the firm invests resources in R&D, the innovation outcome that will result from these R&D investments is yet uncertain.

In a world without private information, the government could perfectly correct for the technology externality through a Pigouvian subsidy and for the non-appropriability of innovation through, for instance, 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 (or with homogeneous firms).

Studying optimal policy under asymmetric information in a dynamic R&D investment model with spillovers is technically involved: we view the tractable model presented in Section 2 to be one of our contributions. In Section 3, we first illustrate the design of optimal policies with asymmetric information in a simplified two-type, one-period toy model. In Section 4, we turn to

the general dynamic, continuous types model, which allows us to crucially study life cycle (age) patterns and the role of the persistence of firm's research productivity. We pose the problem as one of mechanism design, in which we do not 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 for spillovers between agents (firms) in the presence of asymmetric information. With our core setup and methodology in place, additional aspects of R&D investments and innovation by firms can be incorporated, and we discuss possible generalizations and extensions. In Section 5, we solve for and characterize these second-best constrained efficient allocations with asymmetric information and spillovers and highlight the main parameters that determine their sign and magnitude.

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, the formulas are written 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 both 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 to fruitfully address some firm taxation questions which have not been studied until now.

We take the model to the data in Section 6, using firm-level data matched to U.S. Patent Office Patent data, which allows us to see the inputs into R&D, the production decisions, and the innovation output as captured by patents and their citations. We estimate the parameters of our model by matching some key moments of the data, such as the elasticity of the patent quality (measured by citations) to R&D investments, coefficients of variations in patent quality across firms and within firms, growth rates and R&D intensities, or spillover estimates among others. We then numerically simulate and quantify the optimal policies that we previously derived analytically. Finally, we study how close simpler innovation policies can come to approximating the full unrestricted mechanism by comparing the revenue raised from the full optimum to the one raised by restricted and much simpler policies in Section 7.

Our main findings are as follows: Asymmetric information can significantly change the optimal policies. From a theoretical perspective, the constrained efficient incentives for R&D trade-off a Pigouvian correction for the technology spillover and a correction for the monopoly distortion against the need to screen good firms from bad ones. How much R&D should optimally be subsidized depends critically on a key parameter, namely the complementarity of R&D investment to R&D effort (i.e., the complementarity between observable and unobservable innovation inputs) relative to the complementarity of R&D investment to firm research productivity. The more complementary R&D investment is to firm research productivity, the more rents a firm can

extract if R&D investment is subsidized. This puts a brake on how well the government can set the Pigouvian correction and correct for the monopoly distortion. Optimal screening in this case requires dampening the first-best corrective policies. On the other hand, if R&D investments are more complementary to unobservable firm R&D effort, they stimulate the firm to put in more of the unobservable input, which is unambiguously good and would make R&D subsidies optimally larger. Other key determinants of the magnitudes and age-patterns of the optimal policies are the persistence of firms' research productivity shocks and the strength of spillovers. The theory allows us to consider all these different possible patterns; we then turn to the data to see what those important parameters are in reality.

The data shows that R&D investments are highly complementary to firm research productivity: highly productive firms are disproportionately good at transforming R&D inputs into innovation. Given that this implies that higher productivity firms have a comparative advantage at innovation, it is better to incentivize R&D investments less for the lower productivity firms, as this makes mimicking them more attractive for high productivity firms. We discuss how these incentives translate into wedges first and into actual taxes and subsidies second. When it comes to the wedges, a higher net incentive for R&D for higher research productivity firms is provided through the combination of a lower profit wedge at higher profit levels and a lower R&D wedge at higher R&D levels. The intuition for why this is the best screening mechanism is that 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 going to be relatively more attractive to them than it will be to low productivity firms.

Turning to actual taxes and subsidies, we show that it is possible to very closely approximate the optimal unrestricted policies with simple parametric policies. A nonlinear, separable Heathcote-Storesletten-Violante (HSV) type subsidy combined with an HSV-type profit tax reaps almost all of the benefits of the full optimum. It features lower marginal taxes for more profitable firms and lower marginal subsidies at higher R&D investment levels, thus perfectly mirroring the shape of the wedges just described. In fact, we can simplify further, as the most important quantitative feature turns out to be the nonlinearity in the R&D subsidy; making the profit tax linear only generates a small welfare loss. The intuition is that a constant profit tax that is set at a too generous level for low profit firms and at about the right level for high profit firms does reasonably well, since the loss from giving low profit firms a too generous tax is quantitatively small given that they make low profits to start with. Therefore, linear corporate income taxes such as the ones we see around the world can be very close to optimal for innovating firms if combined with the right nonlinear R&D subsidy.

Related Literature. There is a long-standing static contract theory literature on the regulation of firms under private information to which our paper contributes (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.

Our paper, which studies the optimal taxation of firms, is methodologically related to the literature that studies optimal taxation for individuals or households (Saez, 2001). Some recent examples are Kleven (2004), Kleven, Kreiner, and Saez (2009), Lockwood and Weinzierl (2015), Weinzierl (2014), Kindermann and Krueger (2016), and Lockwood (2017). The optimal taxation of firms has been much less explored. Ales and Sleet (2016) examine income taxation when firm productivity is endogenous because of manager talent (but not through the innovation channel as in our paper); Ales, Kurnaz, and Sleet (2015) consider taxation with technical change (i.e., innovation), focusing on an assignment model in the labor market, rather than on the production side (firms) of innovation.

We also contribute to the literature that studies the corrective role of personal income taxes when there are externalities such as rents. In Rothschild and Scheuer (2016), agents can earn income with traditional labor supply, but also engage in rent-seeking. Income taxes will then tend contain some Pigouvian correction for the rent-seeking externality and will vary with the extent to which rent-seeking is at expense of other rent-seekers. Implementing this optimal income tax scheme imposes large informational requirements on the government. In a much simpler model, where it is assumed that the rent-seeking affects everyone in the economy equally, Piketty et al. (2014) derive a simple sufficient statistics formula. Lockwood et al. (2017) allow for both positive-externality occupations and negative-externality ones.⁴ These models are all static, focus on individuals rather than firms, and consider a relatively blunt tool (namely, income taxation) because the externality-inducing action cannot be directly taxed or subsidized. We explicitly model firm taxation and direct subsidies for R&D to incentivize both the observable and unobservable components of it.

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, and where the asymmetric information of the firm regarding 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.

We contribute also 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 (2016), and Werquin (2016). Closest to us in this literature are the papers by Stantcheva (2015) and Stantcheva (2017), which incorporate endogenous investments in human capital into the (personal) dynamic tax problem. In addition to abstracting from any spillovers between agents, this literature also typically merely calibrates the models to simulate optimal policies and

⁴There is also an expansive literature in environmental taxation or behavioral taxation that highlights the correct role of indirect taxes, such as carbon taxes or “sin” taxes.

does not provide a full estimation with a tight link to the data, as we do.

Theoretically, our contributions relative to these aforementioned lines of work are, first, the addition of spillovers between agents (in our case, firms). Because of this important extension, the solution methods are very different, theoretically and computationally. Second, we also provide what is 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 also allow for an infinite horizon. This framework is very malleable: we illustrate several possible extensions in Section 5.2 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 to fully estimate 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.

Even though our analysis is based on the theoretical mechanism design literature, we also use findings from the empirical literature on R&D and productivity to discipline our model and empirical estimation.⁵ First, the empirical evidence on the importance of management practices (Bloom and Van Reenen (2007), Bloom, Sadun, and Van Reenen (2012), Bloom et al. (2013)) lends 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. A large literature documents the important effects of tax incentives for R&D, thus justifying the detailed study of their optimal design. Among many others are the papers by Goolsbee (1998), Bloom, Griffith, and Van Reenen (2002), Bloom and Griffith (2001), and Bloom, Chennells, Griffith, and Van Reenen (2002). In the spirit of the present paper, Serrano-Velarde (2009) examines the heterogeneous impacts of R&D subsidies on firm investments. Syverson (2011) highlights the extent of heterogeneity across firms and surveys the literature trying to understand what causes firms' productivity differences. Taking into account the firm's life cycle – as we do in our model– also seems important given the evidence in Hsieh and Klenow (2014). Finally, Hsieh and Klenow (2009) emphasize the misallocations in factors among firms, which can significantly reduce TFP – the mechanism studied in this paper aims to efficiently allocate innovation factors to heterogenous firms, while taking into account asymmetric information.

Empirically, the overall effects of taxation on innovation or entrepreneurship have been studied by Cullen and Gordon (2006), Cullen and Gordon (2007), Gordon and Lee (2005), Gordon and Sarada (2018), and Akcigit et al. (2018).

Some papers also document the gap between the private and social returns to R&D and, hence, justify the inclusion of non-internalized spillovers into our model and are consistent with

⁵Early papers that pioneered the use of patent data to study firms' innovation choices are Pakes and Griliches (1984), Pakes (1985), and Pakes (1986).

the magnitudes we arrive at in our estimation. Jones and Williams (1998) show that existing econometric estimates of the returns to R&D taken in the context of endogenous growth models imply that optimal R&D investment is at least four times larger than the current one. Jones and Williams (2000) go a step further and try to isolate the channels which determine the net spillover through the lens of a quantitative model. In their endogenous growth model, R&D faces four distortions: it is not appropriable, there are knowledge spillovers, creative destruction and duplication externalities. On net, the authors find that the positive spillovers dominate and there is underinvestment. Bloom, Schankerman, and Van Reenen (2013) also separately estimate business stealing spillovers (through competition in the product market) and knowledge spillovers (in the technology market).

A very large literature studies R&D and innovation policies. A comprehensive, informative, and very recent review of this abundant literature is provided in Bloom et al. (2019). Our main contribution relative to this body of work is to consider asymmetric information about firms' research productivities and to solve for optimal (non-restricted) policies using a mechanism design approach. Grossman et al. (2013) study 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 adopt a Ramsey-approach (linear policies) where they parameterize the policies ex ante and have numerical solutions. We adopt a mechanism design approach. They do not incorporate heterogeneity in and private information about firms' productivities.

While we focus on the design of tax and R&D policies in this paper, there is a large, quite distinct literature on the design of IPR, in which the asymmetric information is typically about the value of the innovation, rather than about the firm's productivity to use research inputs.⁶ A very complementary paper is Chari, Golosov, Tsyvinski (2012) who focus on environments in which the value of the innovation is not known, but there is a market signal about it. When it comes to designing IPR policy, there is also a growing literature on incentivizing innovation through tournaments (Terwiesch and Xu, 2008; Ales et al., 2017). One very interesting finding here is that open tournaments with unrestricted entry are beneficial when the problems are highly uncertain or a large number of possible solutions are needed.

The rest of the paper is organized as follows. Section 2 presents the dynamic model and discusses its assumptions, providing also empirical justification for the focus on asymmetric information in innovation. Section 3 solves a much simpler one-period toy model to draw out some main effects. Section 4 sets up and solves the full dynamic model; Section 5 discusses the forces that shape the optimal policies. Section 6 estimates the model using firm data matched to patent data and simulates the optimal policies. Section 7 considers the welfare loss from simpler,

⁶A few of the many important papers in that literature are Klemperer (1990), Green and Scotchmer (1995), Scotchmer (1999), Gallini and Scotchmer (2002), O'Donoghue and Zweimuller (2004), Hopenhayn et al. (2006), Hopenhayn and Mitchell (2001), Hopenhayn and Mitchell (2011).

restricted policies relative to the full optimum. Section 8 concludes and points to directions for future research.

2 A Dynamic Model of R&D Investments

We present a dynamic model of R&D investments with spillovers that is tractable enough for the theoretical study of the optimal mechanism with asymmetric information. As mentioned in the introduction, first, such a model could lend itself to the study of other types of firm investments with asymmetric information and spillovers. Second, with our core setup and methodology in place, additional aspects of R&D investments by firms can be incorporated. Some of these possible 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 good producers, our framework is also amenable to studying firm taxation with heterogeneous firms more generally, even for non-innovating firms.

2.1 Setting

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:

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

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

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

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 $M''_t(r_t) \geq 0$.⁷ We will 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, among several possible ones, would naturally be unobservable research effort, which is required in order to transform the material resources into an innovation output. We will call these unobservable R&D inputs

⁷Taking 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).

“R&D effort” for concreteness, although they could include other costly, unobservable actions taken at the organizational level and that contribute to research. They 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, the incentive compatibility constraints that will be at the core of our mechanism design problem are needed. In addition, note that if innovation output depends on such unobservable (or non-verifiable) characteristics such as research productivity, it also means that there are some unobservable inputs. The latter prevent the government from perfectly “inverting” the innovation outcome (conditional on observable inputs) to obtain the productivity type. If all inputs were perfectly observable (e.g. 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 , i.e., $\theta^t = \{\theta_1, \dots, \theta_t\}$ and by

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

the probability of that history.

We assume that:

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

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

Let us emphasize here the two related, but conceptually very distinct terms used: Firm’s 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 , that affects the innovation process which produces the product quality q_t .⁸

⁸To clarify a sometimes confusing point: once produced, innovations are non-rival and non-appropriable absent

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 a given 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.⁹ 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, i.e., 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 the 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:

$$\rho_{xy} = \frac{\frac{\partial^2 \lambda}{\partial x \partial y} \lambda}{\frac{\partial \lambda}{\partial x} \frac{\partial \lambda}{\partial y}} \quad (3)$$

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:

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

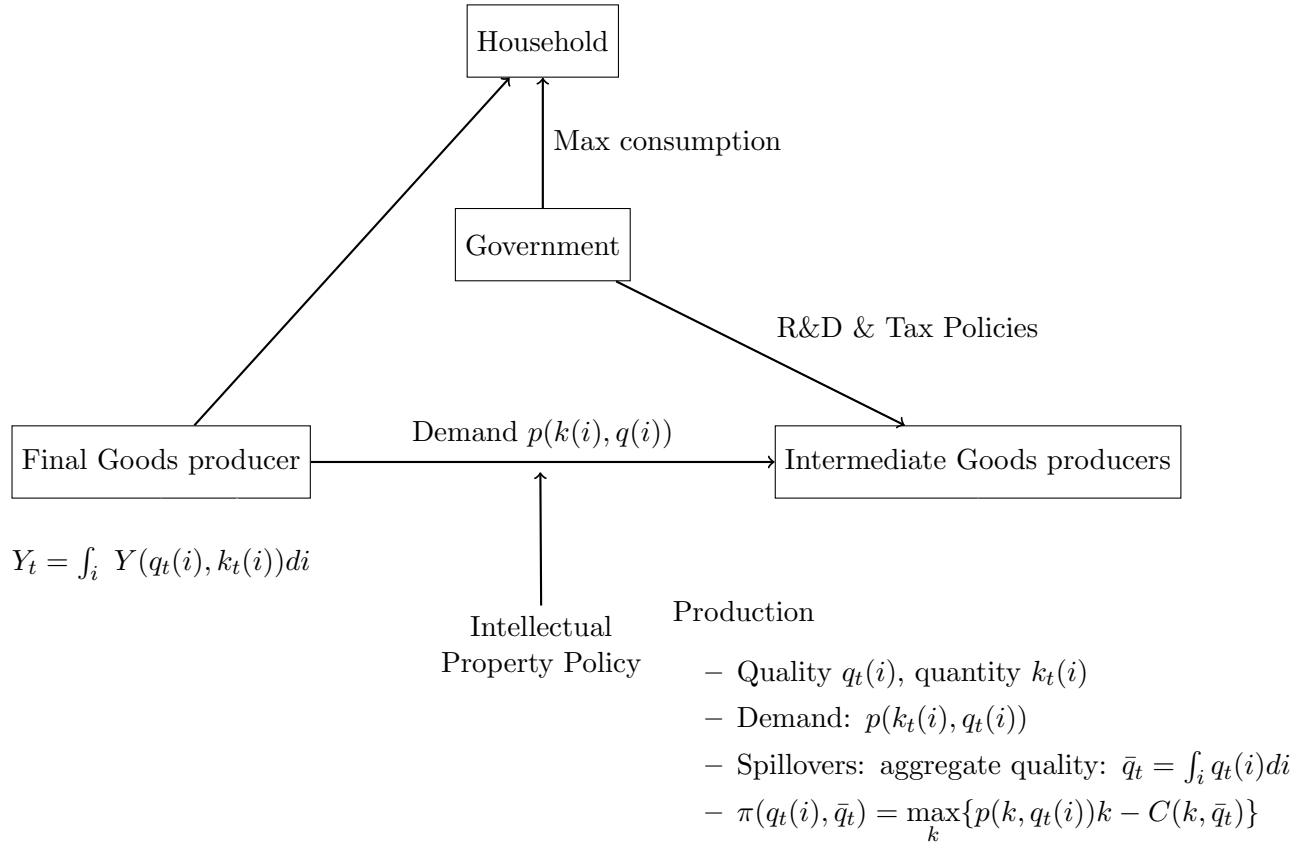
has $\rho_{\theta l} = \rho_{\theta r} = \rho_{lr} = \rho_t$.

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 learn-

IPR. The inputs into that innovation are, as usual, rival.

⁹If the timing was contemporaneous, the sums and expectations for the R&D wedge in Proposition 2 should simply be lagged one period.

FIGURE 1: MODEL SUMMARY



ing from new technologies to improve one's production. Currently, 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.¹⁰ Aggregate quality is given by:

$$\bar{q}_t = \int_{\Theta^t} q_t(\theta^t) P(\theta^t) d\theta^t \quad (4)$$

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

Final goods production: The final good is consumed by consumers and is produced competi-

¹⁰In 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.

tively using the intermediate goods as inputs. The production technology for the final good is:

$$Y_t = \int_{\Theta^t} Y(q_t(\theta^t), k_t(\theta^t)) P(\theta^t) d(\theta^t) \quad (5)$$

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 one. The demand function for the intermediate good that arises in the market will depend on the IPR regime, to which we turn next.

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 good producer has monopoly power and faces a downward sloping demand curve derived from the optimization problem of the final good 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 T periods, which in our benchmark will be an infinite number of periods ($T = \infty$). It is, however, possible to accommodate a finite life cycle in our framework, by simply replacing the infinite forward-sums with forward-sums up to some $T < \infty$. 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 :

$$\sum_{t=1}^{\infty} \left(\frac{1}{R}\right)^{t-1} \int_{\Theta^t} (p(q_t(\theta^t), k_t(\theta^t)) k_t(\theta^t) - C(k_t(\theta^t), \bar{q}_t) - M_t(r_t(\theta^t)) - \phi_t(l_t(\theta^t))) P(\theta^t | \theta_1) d(\theta^t | \theta_1) \quad (6)$$

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:

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

Their maximization pins down the quantity produced for a given quality level. Figure 1 summarizes the model in schematic 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, the gross transfer is just price times quantity and the firm payoff is as in (6)). 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 , equal to:

$$v_t(\theta^t) = T_t(\theta^t) - \phi_t(l_t(\theta^t)) \quad (7)$$

Social welfare (the objective the planner maximizes) is a weighted sum of consumer surplus plus firm surplus:¹¹

$$\sum_{t=1}^{\infty} \left(\frac{1}{R}\right)^{t-1} \left(\int_{\Theta^t} (Y(k_t(\theta^t), q_t(\theta^t)) - (C(k_t(\theta^t), \bar{q}_t) + M_t(r_t(\theta^t)) + T_t(\theta^t)) + (1 - \chi)v_t(\theta^t)) P(\theta^t) d(\theta^t) \right) \quad (8)$$

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 such 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 one (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 ($\chi = 0$).¹² 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:

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

with q_0 and r_0 given.

¹¹The final goods producer always has zero payoff because it operates under perfect competition.

¹²Under full information, type-specific lump-sum transfers and taxes are feasible.

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

$$q_t = (1 - \delta)q_{t-1} + \lambda_t \quad \text{with} \quad 0 < \delta < 1 \quad (9)$$

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

The optimal choice of R&D investment and firm effort is then such 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 of 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 for instance 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 good producer and has to take as given their production decisions. In section 5.4, 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 5.4. 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 5.3). We subsequently study the shape of and revenue losses from restricted, parametric instruments, which are simpler (Section 7).

The Importance of Asymmetric Information: We now highlight why asymmetric information is a critically important feature in the innovation process. We do this in two ways. First, we summarize the abundant literature showing the prevalence of asymmetric information; second, we show directly 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, of its business model, or of its 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.¹³ A large literature indeed argues that asymmetric information is likely to be a key issue in innovation. Hall and Lerner (2009) summarize several of these contributions. In their language, 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 – that leads to financing frictions and inefficiencies– they argue 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 caution 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 very much 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 likely, such as high-technology and innovating sectors. Kaplan and Stromberg (2001) also document 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 et al., 2014). Kerr et al. (2014) use data from Thompson Venture Economics and calculate the returns made by VCs on every startup between 1985 and 2009. More than 55% of all startups that received VC funding terminated at a loss; however a small share of 6% of all startups generated more than 500% return. The returns are thus very skewed and even experienced investors are not able to screen firms perfectly. This is also showcased by the following results: asked to assign scores to innovators, the distribution of scores by expert investors for investments that turned

¹³Recall 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.

out to be successful was statistically not different from the distribution of scores assigned to investments that turned out to be unsuccessful. To mitigate the asymmetric information problems, VCs tend to focus on specific sectors and provide staged investments, where subsequent funding rounds are conditional on success in earlier stages. They thus screen firms over time, as their productivity and success potential is revealed. Again, given that VCs face asymmetric information problems despite the huge time investment and detailed involvement in the innovating 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.

A strand of the literature has looked at stock market responses as another symptom of the asymmetric information problem inherent in innovation. Event studies that measure the market response to announcement of new debt or share issue (considered good news) show that there are higher abnormal returns to firm shares when the firm is more R&D-intensive (Zantout, 1997; Alam and Walton, 1995). In addition, Gharbi et al. (2014) highlight, using French high-tech firms, that stock returns are more volatile for more R&D-intensive firms because the information asymmetry generated by innovation makes prospects more uncertain. Aboody and Lev (2000) show that insider gains are larger at R&D intensive firms than firms without R&D because that is where asymmetric information is higher. Insiders also appear to take advantage of planned changes in R&D budgets. They conclude that R&D is a major contributor to information asymmetry that ultimately leads to insider gains.

Finally, the literature also makes it very clear that not all R&D inputs are easily verifiable. Hall and Van Reenen (2000) call this the “relabeling” problem and offer many examples of it. Mansfield (1986) surveys the effects of R&D tax credits in the US, Canada, and Sweden and finds that there is substantial misreporting. More recently, Chen et al. (2019) document 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. To do so, we study what share of the innovation quality of a firm can actually be predicted based on its past and current observables. The data and measurement are explained in more detail in Section 6. 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 et al., 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 result is that the prediction is quite poor. Even adding such an exhaustive list of control variables – which already goes well beyond what a government could conceivably condition a tax system on – 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” i.e., 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. We want to caution, however, that this is not a “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). However, this analysis does suggest that it is very difficult to predict the quality of a firm’s innovation, which creates a particularly fertile ground for asymmetric information.

Three properties of the firm’s stochastic research productivity process matter for the optimal policies, as will be clear in the formulas in Section 5: its initial dispersion (heterogeneity), its variance, and its persistence. To make sure that the stochastic properties in the model match those of the data, we estimate these parameters directly in the data based on the residual within-firm and across-firm dispersions in quality (controlling for observables). We also provide many robustness checks specifically on the stochastic process in Section 6.5.

2.5 Discussion of the Assumptions and Possible Generalizations

1. **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 finetune 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 discuss heterogeneity in firms’ production productivities explicitly in Section 5.2.
2. **Entry and Exit:** In principle, firms in our model only make intensive margin decisions about how much to produce. Exit and entry can, however, be captured to some extent. Regarding exit, the corner solution of zero production could represent exit. Empirically, firms at different ages have heterogeneous exit rates, with exit rates declining with age conditional on having survived until that age. This can be captured by letting R depend on firm 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.¹⁴
3. **The role of IPR:** Our focus is not on IPR, but on the design of R&D policies. However, the latter’s shape and magnitudes do 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 hence,

¹⁴For instance, the government could endogenously set the lower bound of θ that would optimally be allowed to enter.

full monopoly rights to innovating firms.¹⁵ As a result, part of the role of R&D policies will be to correct partially for the monopoly distortion induced by the patent system.¹⁶ We also consider two different cases based on whether the government can intervene in the private market between intermediate and final good producers, i.e., whether it can observe and make the optimal policy contingent on the quantity k produced. Whenever the government can control quantity, given that quality is observable, it can incentivize the socially optimal quantity to be produced and thus counteract the monopoly distortion (see Section 5.4).

4. **Shape of the spillovers:** The exact shape of the spillovers will not be important for the theoretical results and will not affect the forces we describe and the key qualitative 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 the 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 own quality and aggregate quality. This is a quite general formulation that will apply for many specifications of where the spillover applies and regardless of 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.

5. **Different types of investments with different externalities:** It is possible to consider different types of firm investments, with different externalities. Section 5.2 studies these cases in detail.

¹⁵It would be easy to instead consider another system, such as patent for protection for x years or patent protection for a fraction of the monopoly profits.

¹⁶If the world were different, and there was an IPR policy that did not grant monopoly power, e.g.: 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 good 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.

3 Optimal Policies in a Simple Two-type, One-Period Model

In this section, we illustrate the underlying logic of the optimal mechanism in a very simple two-type, one-period model.

Suppose that firms can be of the high research productivity type θ_2 or of the low productivity type θ_1 . The fractions in the population of firms of types high and low are, respectively, f_2 and f_1 , with $f_2 = 1 - f_1$. The problem is static: Firms enter period 1 with a knowledge of their type realization, chose R&D investments $r(\theta_i)$ and R&D effort $l(\theta_i)$ at the beginning of the period. The step size is $\lambda(\theta_i) = \lambda(r(\theta_i), l(\theta_i), \theta_i)$ and quality is $q(\theta_i) = q_0 + \lambda(\theta_i)$, where q_0 is given. At the end of the period firms receive a transfer $T(\theta_i)$ from the government. For the exposition, suppose that the step size takes the form:

$$\lambda(r, l, \theta_i) = w(r, \theta_i)l$$

for an increasing and concave function w . Profits are denoted by $\pi(q, \bar{q})$ as a function of quality q and aggregate quality $\bar{q} = f_1q(\theta_1) + f_2q(\theta_2)$.

In the planning problem, the planner sets a menu of contracts $(r(\theta_i), l(\theta_i), T(\theta_i))$ for $i = 1, 2$ and lets firms self-select allocations from this menu. For simplicity, we set $\chi = 1$.¹⁷ For any quality, the firm will choose the privately optimal quantity, leading to output net of production costs $\tilde{Y}(q(\theta_i), \bar{q})$ for type θ_i . The remaining components of the menu $(r(\theta_i), l(\theta_i), T(\theta_i))_{i=1,2}$ and \bar{q} are chosen to maximize social welfare defined in (8), and which in this simple case becomes:

$$W = f_1 (\tilde{Y}(q(\theta_1), \bar{q}) - M(r(\theta_1)) - T(\theta_1)) + f_2 (\tilde{Y}(q(\theta_2), \bar{q}) - M(r(\theta_2)) - T(\theta_2))$$

subject to $q(\theta_i) = q_0 + \lambda(\theta_i)$ with q_0 given, and subject to firms' participation constraints:

$$T(\theta_i) - \phi(l(\theta_i)) \geq 0$$

We can also allow for some different thresholds in the participation constraint, such that $T(\theta_i) - \phi(l(\theta_i)) \geq \underline{V}(\theta_i)$. In the first best, firm type is observable, $\chi = 0$, and the planner makes each firm invest the efficient level of effort and inputs, such that the marginal effort and R&D investment costs equal the social impact, as in section 2.3, and surplus is extracted in a lump-sum fashion from the firms, i.e.,¹⁸

$$T(\theta_i) = \phi(l(\theta_i))$$

¹⁷This is without loss of generality: a $\chi \neq 1$ would simply appear as a scaling factor in front of the screening term in the formulas below.

¹⁸More precisely,

$$M'(r(\theta_i)) = \left(\frac{\partial \tilde{Y}(q(\theta_i), \bar{q})}{\partial q} + \left(f_1 \frac{\partial \tilde{Y}(q(\theta_1), \bar{q})}{\partial \bar{q}} + f_2 \frac{\partial \tilde{Y}(q(\theta_2), \bar{q})}{\partial \bar{q}} \right) \right) \frac{\partial \lambda(r(\theta_i), l(\theta_i), \theta_i)}{\partial r(\theta_i)}$$

$$\frac{\phi(l(\theta_i))}{w(r(\theta_i), \theta_i)} = \frac{\partial \tilde{Y}(q(\theta_i), \bar{q})}{\partial q} + \left(f_1 \frac{\partial \tilde{Y}(q(\theta_1), \bar{q})}{\partial \bar{q}} + f_2 \frac{\partial \tilde{Y}^*(q(\theta_2), \bar{q})}{\partial \bar{q}} \right)$$

The second-best problem imposes an incentive constraint for each type i :

$$T(\theta_i) - \phi(l(\theta_i)) \geq T(\theta_j) - \phi\left(\frac{w(r(\theta_j), \theta_j)l(\theta_j)}{w(r(\theta_j), \theta_i)}\right) \quad \forall (i, j)$$

Given that the goal is to minimize total transfers to the firms, one can show that the incentive constraint of type θ_2 and the participation constraint of type θ_1 will be binding.¹⁹ Indeed, at the first-best allocations and transfer levels, high research productivity firms will be tempted to pretend that they are low productivity firms. This is because they have to forfeit all their surplus to the planner, but, since they are able to reach any step size at a lower R&D effort cost than low research productivity firms, they could achieve a positive surplus by selecting the low research productivity firm's first-best allocation. To prevent this from happening, the allocation of the low research productivity firms needs to be distorted so as to make it less attractive to high productivity firms.

The transfers then have to satisfy:

$$T(\theta_1) = \phi(l(\theta_1))$$

$$T(\theta_2) - \phi(l(\theta_2)) \geq T(\theta_1) - \phi\left(\frac{w(r(\theta_1), \theta_1)l(\theta_1)}{w(r(\theta_1), \theta_2)}\right)$$

Substituting these expressions into the social objective, we obtain the so-called virtual surplus, which is social surplus minus the informational rent forfeited to the high type θ_2 to induce him to truthfully reveal his type. The social optimum will maximize allocative efficiency (the first line below) while trying to reduce the informational rent forfeited to the high type (the second line).

$$W = f_1 (\tilde{Y}(q_1(\theta_1), \bar{q}) - M(r(\theta_1)) - \phi(l(\theta_1))) + f_2 (\tilde{Y}(q(\theta_2), \bar{q}) - M(r(\theta_2)) - \phi(l(\theta_2))) - f_2 \left(\phi(l(\theta_1)) - \phi\left(\frac{w(r(\theta_1), \theta_1)l(\theta_1)}{w(r(\theta_1), \theta_2)}\right) \right) \quad (10)$$

Characterization of the optimal allocation in terms of wedges. The constrained efficient allocation can be described using so-called wedges or implicit taxes and subsidies, which measure the deviation of the allocation relative to the laissez-faire economy with patent protection. In the laissez-faire economy with patent protection, profits are a function of the product's quality and aggregate quality, $\pi(q(\theta_i), \bar{q})$, as defined in Section 2. The effort wedge, $\tau(\theta_i)$ on type θ_i is defined as the gap between the marginal *private* benefit of effort and its cost, while the R&D investment wedge is defined as the gap between the marginal cost of R&D and its marginal private benefit. Thus, a higher effort wedge means a lower incentive for R&D effort, while a higher R&D

¹⁹As is usual in these types of screening problems, the slackness of the low type's omitted incentive constraint can be checked *ex post*.

investment wedge means a higher incentive for R&D investments. Formally:

$$s(\theta_i) = M'(r(\theta_i)) - \frac{\partial \pi(q(\theta_i), \bar{q})}{\partial q(\theta_i)} \frac{\partial \lambda(r(\theta_i), l(\theta_i), \theta_i)}{\partial r(\theta_i)}$$

$$(1 - \tau(\theta_i)) \frac{\partial \pi(q(\theta_i), \bar{q})}{\partial q(\theta_i)} \frac{\partial \lambda(r(\theta_i), l(\theta_i), \theta_i)}{\partial l(\theta_i)} = \phi'(l(\theta_i))$$

In the implementation below, it will be clear that there is a very natural map between the wedges (i.e., implicit taxes and subsidies) and the explicit marginal tax rates of the implementing tax function.

Taking the first-order conditions of the social objective with respect to $r(\theta_i)$ and $l(\theta_i)$ for $i = 1, 2$ and using the definitions of the wedges, we obtain that for the low research productivity type, the allocations are distorted just enough to balance the informational rent forfeited to the high type and the loss in allocative efficiency.

Proposition 1. Optimal Allocations for Low Research Productivity Firms.

i) The optimal R&D investment wedge on the low research productivity type is given by:

$$s(\theta_1) = \underbrace{\left(f_1 \frac{\partial \tilde{Y}(q(\theta_1), \bar{q})}{\partial \bar{q}} + f_2 \frac{\partial \tilde{Y}(q(\theta_2), \bar{q})}{\partial \bar{q}} \right) \frac{\partial w(r(\theta_1), \theta_1)}{\partial r} l(\theta_1)}_{\text{Pigouvian correction}} + \underbrace{\left(\frac{\partial \tilde{Y}(q(\theta_1), \bar{q})}{\partial q(\theta_1)} - \frac{\partial \pi(q(\theta_1), \bar{q})}{\partial q(\theta_1)} \right) \frac{\partial w(r(\theta_1), \theta_1)}{\partial r} l(\theta_1)}_{\text{Monopoly quality valuation correction}}$$

$$+ \underbrace{\frac{f_2}{f_1} \left(1 - \frac{\frac{\partial \log(w(r(\theta_1), \theta_2))}{\partial \log(r)}}{\frac{\partial \log(w(r(\theta_1), \theta_1))}{\partial \log(r)}} \right) \frac{\partial w(r(\theta_1), \theta_1)}{\partial r} l(\theta_1)}_{\text{Complementarity}} \phi' \left(\frac{w(r(\theta_1), \theta_1) l(\theta_1)}{w(r(\theta_1), \theta_2)} \right)}_{\text{Screening term}} \quad (11)$$

ii) The optimal R&D effort wedge on the low productivity firm is given by:

$$\tau(\theta_1) \frac{\partial \pi(q(\theta_1), \bar{q})}{\partial q(\theta_1)} = - \left(\frac{\partial \tilde{Y}(q(\theta_1), \bar{q})}{\partial q(\theta_1)} - \frac{\partial \pi(q(\theta_1), \bar{q})}{\partial q(\theta_1)} \right) - \left(f_1 \frac{\partial \tilde{Y}(q(\theta_1), \bar{q})}{\partial \bar{q}} + f_2 \frac{\partial \tilde{Y}(q(\theta_2), \bar{q})}{\partial \bar{q}} \right)$$

$$+ \frac{f_2}{f_1} \left(\frac{1}{w(r(\theta_1), \theta_1)} \phi'(l(\theta_1)) - \frac{1}{w(r(\theta_1), \theta_2)} \phi' \left(\frac{w(r(\theta_1), \theta_1) l(\theta_1)}{w(r(\theta_1), \theta_2)} \right) \right) \quad (12)$$

Screening term: Cost differential between high and low productivity firms

Proof. See Appendix A.2. □

The optimal implicit subsidy on R&D investment in (11) and the R&D effort wedge in (12) balance three considerations.

1) *Pigouvian correction for technology spillovers:* Incentives are increasing in the Pigouvian correction that aligns private incentives with the social benefit from R&D technology spillovers, which are the key reason for the government to intervene. This correction is larger when the marginal return to R&D investments ($\frac{\partial w(r(\theta_1), \theta_1)}{\partial r}$) is larger.

2) *Monopoly quality valuation correction*: Starting from a laissez-faire with patent protection, the monopolist values each marginal increase in quality less than its marginal social value: this difference in quality valuation must also be corrected for in the optimal planning problem. This is the second term in each of the wedge formulas. The distortions in the R&D investment and effort are modified so as to indirectly compensate for the under-provision of quantity of the monopolist. The effect of a change in quantity (induced by extra investment in R&D investment or R&D effort) on social welfare, implicit in $\frac{\partial \tilde{Y}(q(\theta_i), \bar{q})}{\partial q(\theta)}$, is first-order and is proportional to the monopoly distortion, i.e., the gap between price and marginal cost.²⁰ The pre-existing monopoly distortions amplifies the direct impact of R&D effort and investment on output and the indirect impact through the technology spillover, pushing the R&D effort wedge down and the R&D investment wedge up. The optimal R&D policies hence depend on the IPR policies in place. If there was no monopoly distortion in the laissez-faire economy, i.e., if there was for instance a prize system, then there would be no need to correct for it and this term would disappear from the optimal wedge formulas.²¹

3) *Screening term*: The screening term (the third term in each formula) captures the modification to the first-best incentive that is induced by the asymmetric information. It is decreasing in the fraction of high research productivity firms over low research productivity firms: the lower the fraction of low productivity firms, and the less costly it is to distort their effort or investments for the sake of reducing the informational rent of the (more frequently encountered) high productivity firms.

The screening term depends on the relative complementarity of R&D investments with R&D effort versus with firm research productivity. Since the step size is assumed here to be multiplicatively separable, the elasticity of the step size to R&D effort for both types is just 1, the first term in the “complementarity” term. The relative elasticity of the return to effort $w(r, \theta)$ with respect to R&D for the high and the low type, $\frac{\partial \log(w(r(\theta_1), \theta_2))}{\partial \log(r)} / \frac{\partial \log(w(r(\theta_1), \theta_1))}{\partial \log(r)}$ measures how complementary R&D investments are to firm research productivity: if the elasticity is increasing in type, then R&D investments benefit disproportionately high research productivity firms. The more elastic the high type’s return is to R&D, the less the R&D investment of the low type can be subsidized, as this makes it more tempting for the high type to pretend to be low type. Put differently, increasing R&D investments of the low type when the relative elasticity is high means tightening the high type’s incentive constraint and giving that firm more informational rent. As a special case, if the elasticities of the high and low types are the same, then R&D investments of the low type do not affect the high type’s incentive constraint. As a result, the screening term drops out and the optimal marginal R&D subsidy is set solely to correct for the technology spillover and the monopoly distortion.

²⁰Formally, $\frac{\partial \tilde{Y}(q(\theta_i), \bar{q})}{\partial q(\theta_i)} = \frac{\partial Y(q(\theta_i), k(q(\theta_i), \bar{q}))}{\partial q(\theta_i)} + \left(p(q(\theta_i), k(q(\theta_i), \bar{q})) - \frac{\partial C(k(q(\theta_i), \bar{q}), \bar{q})}{\partial k} \right) \frac{\partial k(q(\theta_i), \bar{q})}{\partial q(\theta_i)}$ where $k(q(\theta_i), \bar{q})$ is the quantity chosen to maximize profits by a monopolist with quality $q(\theta_i)$.

²¹Naturally, larger wedges (i.e., distortions relative to the laissez-faire) do not imply in any sense that there is more investment in effort or R&D relative to a situation with smaller wedges.

Stimulating R&D investments is beneficial when there is a high complementarity of R&D investments with unobservable R&D effort, because it stimulates the unobservable input, but is detrimental when there is a high complementarity with firm research productivity, as it then tightens the incentive constraint of the high research productivity firm. The basic logic is that investments in R&D are distorted only in so far as they (beneficially) affect the incentive constraint of the high research productivity firm, i.e., as long as they can indirectly stimulate the unobservable R&D effort choice.

For the R&D effort wedge, the efficiency cost of distorting the low research productivity firm's R&D efforts depends on the comparative productive advantage on the high type relative to the low type. The efficiency cost depends on the difference in the marginal cost $\phi'(l)$ of producing the step size assigned to the low research productivity firm (which is $\lambda(\theta_1)$) between the low and the high research productivity firm. Since the cost function $\phi(l)$ is convex, this difference is always positive. The smaller this difference the more tempting it is for the high research productivity firm to imitate the low research productivity one and the more the R&D effort of low productivity firms should be reduced. This increases the optimal effort wedge $\tau(\theta_1)$ on the low productivity firm's R&D effort.

On the other hand, the high research productivity firms' allocations are set based on the monopoly valuation and Pigouvian correction terms only. The screening term is zero since the low type's incentive constraint is not binding. Appendix A.1 explains two possible implementations of the optimal allocations in this simple model and provides expressions for the marginal tax rates and the marginal subsidy rate in the case in which this implementing tax system can be made differentiable.

4 A Dynamic Direct Revelation Mechanism with Spillovers

We now return to the general dynamic model from Section 2. Recall that the history θ^t and research effort l_t are private information of each firm. 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 $r(\theta^0) = r_0$.²² 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

²²Since 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 fundamentally, but makes the notation heavier.

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:

$$V^\sigma(\theta^t) = T_t(\theta^t(\theta^t)) - \phi_t(l_t(\lambda_t(\theta^t(\theta^t))), r(\theta^{t-1}(\theta^{t-1}), \theta_t)) + \frac{1}{R} \int_{\Theta} V^\sigma(\theta^{t+1}) f^{t+1}(\theta_{t+1}|\theta_t) d\theta_{t+1}$$

$V^\sigma(\theta^t)$ 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:

$$V_1(\{\lambda(\theta^s), r(\theta^s), T_s(\theta^s)\}_{s=1}^{\infty}, \theta_1) = \sum_{t=1}^{\infty} \left(\frac{1}{R}\right)^{t-1} \cdot \left\{ \int_{\Theta^t} \{T_t(\theta^t) - \phi_t(l_t(\theta^t))\} P(\theta^t|\theta_1) d\theta^t \right\} \\ \text{with } l_t(\theta^t) := l_t(\lambda_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t) \quad (13)$$

4.1 A first-order approach

We use a first-order approach, which replaces all the incentive constraints of agents by their envelope conditions.²³ 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 by optimality of the report).

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:

$$I_{1,t}(\theta^t) = \prod_{s=2}^t \left(-\frac{\frac{\partial F^s(\theta_s|\theta_{s-1})}{\partial \theta_{s-1}}}{f^s(\theta_s|\theta_{s-1})} \right) \quad (14)$$

The impulse response function captures the persistence of the stochastic type process. For instance, for an autoregressive process where $\theta_t = \bar{p}\theta_{t-1} + \varepsilon_t$, the impulse response is simply $I_{1,t}(\theta^t) = \bar{p}^{t-1}$. 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).

²³See Pavan et al. (2014), Farhi and Werning (2013), and Stantcheva (2017).

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

This is the full support assumption that can be relaxed as in Farhi and Werning (2013) to allow for 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 period's $s + 1$ 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. The assumption of boundedness could instead be replaced by 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), iid, or a fully persistent process satisfy this assumption.

We can rewrite the per-period payoff of the firm from (7) as a function of the allocation of transfer, step size, and past R&D spending and given its true type θ_t :

$$v_t(T_t, \lambda_t, r_{t-1}; \theta_t) = T_t - \phi_t(l_t(\lambda_t, r_{t-1}, \theta_t)) \quad (15)$$

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

$$\frac{\partial V(\theta^t)}{\partial \theta_t} = \mathbb{E} \left(\sum_{s=t}^{\infty} I_{t,s}(\theta^s) \left(\frac{1}{R} \right)^{s-t} \frac{\partial v_s(\theta^s)}{\partial \theta_s} \mid \theta_t \right) \quad (16)$$

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 :

$$V_1(\theta_1) \geq 0 \quad (17)$$

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

$$V(\theta^{t-1}, \theta_t) = \int_{\underline{\theta}}^{\theta_t} \frac{\partial V(\theta^{t-1}, m)}{\partial m} dm + V(\theta^{t-1}, \underline{\theta}) \quad (18)$$

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

4.2 Planner's problem

The planner's objective is to maximize social welfare in (8) subject to the incentive constraints in (16) and participation constraints in (17). For simplicity only, we set $\chi = 1$.²⁴

²⁴As explained, above, this 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.

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 good 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 (13), 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 $V_1(\theta_1) = 0$. Using the expression for the informational rent that needs to be forfeited to each agent from (18), the expected discounted payoff to the planner is the “virtual surplus,” i.e., the total social surplus minus informational rents.

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

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 optimal allocations subject to resource and incentives constraints are solved for as functions of this conjectured sequence. To ensure that the sum of aggregate qualities that arises is consistent with the conjectured \bar{q} , a consistency constraint needs to be imposed for every period t :

$$\int_{\Theta^t} q_t(\theta^t) P(\theta^t) d\theta^t = \bar{q}_t \quad (19)$$

Let η_t 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:

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

$$\int_{\Theta^t} q_t(\theta^t) P(\theta^t) d\theta^t = \bar{q}_t \quad \text{and} \quad q_t(\theta^t) = q_{t-1}(\theta^{t-1})(1 - \delta) + \lambda(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t) \quad (20)$$

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

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

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.

$$P : \max_{\bar{q}} P(\bar{q}) \tag{21}$$

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, i.e., that the global incentive constraints are satisfied.²⁵ This numerical verification procedure is described in Appendix A.2. For the range of parameters we study in Section 6, the allocations found using the second-order approach do indeed satisfy the global incentive constraints.²⁶ In addition, the optimal allocations in such dynamic models with spillovers can not 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 6.

4.3 Characterizing the Constrained Efficient Allocation Using Wedges

To characterize the constrained efficient allocations it is very helpful, as in the simple illustration of Section 3, 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, i.e., the hypothetical incentives expressed as implicit taxes or subsidies that would have to be provided to firms starting from the laissez-faire 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 in the laissez-faire with patent protection.

²⁵See also Farhi and Werning (2013) and Stantcheva (2017).

²⁶All code files will be made available to interested researchers.

Definition 1. *The corporate wedge and the R&D wedge.* The corporate (or profit) wedge is defined as:

$$\tau(\theta^t) := \mathbb{E} \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \frac{\partial \lambda_t(\theta^t)}{\partial l_t(\theta^t)} \right) - \phi'(l_t(\theta^t)) \right) \quad (22)$$

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

$$s(\theta^t) := M'_t(r_t(\theta^t)) - \frac{1}{R} \mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \frac{\partial \lambda_{t+1}(\theta^{t+1})}{\partial r_t(\theta^t)} \right) \right) \quad (23)$$

To simplify the notation, we use the following definitions.

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

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

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

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

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

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.²⁷

5 Optimal Policies

In this section, we characterize the optimal constrained efficient allocations that are the solutions to the planning problem in Section 4. We then show how these allocations can be implemented with a parsimonious tax function.

5.1 Optimal Corporate and R&D Wedges

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

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

²⁷Note that since the quantity maximizes consumption net of production costs per producer, i.e., 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).

For instance, $\varepsilon_{l(1-\tau),t}$ is the elasticity of R&D effort to the net-of-tax rate $1 - \tau$. 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 2. Optimal corporate wedge and R&D wedge.

(i) The optimal profit wedge satisfies:

$$\tau(\theta^t) = \underbrace{-\mathbb{E} \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \eta_s \right) \frac{\partial \lambda_t}{\partial l_t}}_{\text{Pigouvian correction}} - \underbrace{\mathbb{E}(Q_t(\theta^t) - \Pi_t(\theta^t)) \frac{\partial \lambda_t}{\partial l_t}}_{\text{Monopoly quality valuation correction}} + \underbrace{\frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t}(\theta^t)}_{\text{Type distribution and persistence}} \underbrace{\frac{\phi'_t \lambda_{\theta t}}{\lambda_t} \left[\frac{1}{\varepsilon_{l,1-\tau}} \frac{1}{\varepsilon_{\lambda l,t}} + \rho_{\theta l,t} \right]}_{\text{Screening and incentive term Elasticity}} \quad (24)$$

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

$$s(\theta^t) = \underbrace{\mathbb{E} \left(\sum_{s=t+1}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t-1} \eta_s \frac{\partial \lambda(\theta^{t+1})}{\partial r_t} \right)}_{\text{Pigouvian correction}} + \underbrace{\mathbb{E} \left((Q_{t+1}(\theta^{t+1}) - \Pi_{t+1}(\theta^{t+1})) \frac{\partial \lambda(\theta^{t+1})}{\partial r_t} \right)}_{\text{Monopoly quality valuation correction}} \quad (25)$$

$$+ \frac{1}{R} \mathbb{E} \left(\underbrace{\frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t+1}(\theta^{t+1}) \phi'_{t+1}(l(\theta^{t+1}))}_{\text{Type distribution and persistence}} \frac{\lambda_{\theta} \lambda_r}{\lambda \lambda_l} \underbrace{(\rho_{lr} - \rho_{\theta r})}_{\text{Relative complementarity}} \right)$$

Screening and incentive term

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

$$\int_{\Theta^t} \frac{\partial \tilde{Y}^*(q_t(\theta^t), \bar{q}_t)}{\partial \bar{q}_t} P(\theta^t) d\theta^t = \eta_t \quad (26)$$

Proof. See Appendix A.2. □

The optimal wedges in (24) and (25) 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 good 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, i.e., the gap between price and marginal cost, cumulated over all future periods.²⁸ 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 higher the gap in between the monopolist's 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's 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 strength of share of total value granted to innovating firms by the IPR system. The closer the company is to capturing the full social surplus and 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 pushes towards increasing firms' R&D effort and investment relative to the laissez-faire. The Pigouvian correction for R&D effort in (24) 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 (25) is increasing in the expected effect of R&D investments on the next period's step size $\frac{\partial \lambda_{t+1}}{\partial r_t}$.

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 just described monopoly and Pigouvian corrections. The screening term arises because of the 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^t), and there would be no informational rents. With asymmetric information, on the contrary, there are three effects at play.

²⁸Formally, $\frac{\partial \tilde{Y}(q_t(\theta^t), \bar{q}_t)}{\partial q} = \frac{\partial Y(q_t(\theta^t), k_t(q_t(\theta^t), \bar{q}_t))}{\partial q_t(\theta^t)} + (p(q_t(\theta^t), k_t(q_t(\theta^t), \bar{q}_t)) - \frac{\partial C(k_t(q_t(\theta^t), \bar{q}_t), \bar{q}_t)}{\partial k}) \frac{\partial k_t(q_t(\theta^t), \bar{q}_t)}{\partial q_t(\theta^t)})$ 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.

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 iid, 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.

The higher the inverse hazard ratio $\frac{1-F^1(\theta_1)}{f^1(\theta_1)\theta_1}$ and the larger the mass of firms with research productivity larger than θ_1 relative to the mass of firms with type θ_1 ($f^1(\theta_1)$). This 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.²⁹ The efficiency cost can be decomposed in allocative inefficiency and information rents. The allocative inefficiency induced by the effort wedge is increasing in the elasticity of the step size 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 $\rho_{\theta l,t}$. Recall from the simple illustration in Section 3 that the effort wedge on the low research productivity firm was higher when the high research productivity firm was able to mimic its step size production at a much lower effort cost. This is the effect embodied in $\rho_{\theta l}$. 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 those rents. Since the disutility of R&D effort is indexed by t , the strength of this incentive effect could be varying over the life cycle of a firm.

The complementarity between R&D, firm effort, and firm type. Recall from Section 3 that for the purposes of screening, the (observable) R&D investments are distorted only in so far as they can indirectly affect the unobservable R&D effort choice, i.e., can affect the incentive constraint of the high research productivity firm.

How effective R&D investment subsidies are to stimulate 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 the firm as long as they increase the marginal return to effort, i.e., 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

²⁹This is naturally reminiscent of the inverse elasticity rule in Ramsey taxation.

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 forfeited 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, i.e., 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 $(\rho_{lr} - \rho_{\theta r}) > 0$, i.e., 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.³⁰ The higher this complementarity, and 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. It could vary with the level of effort, R&D, and ability, as well as with firm age.³¹ Hence, the optimal R&D wedge may change sign over the distribution of types or over the life cycle of a firm.

5.1.1 Cross-sectional profile of optimal policies

Given the level of generality at which we have derived the wedges, there are many possible patterns by firm productivity. But we can discuss, term-by-term, what forces would drive these cross-sectional patterns.

When thinking of the latter, 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 the *laissez-faire*. This is because firms have heterogeneous benefits and costs from investments and effort in the *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 in general 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, the latter's allocations will be distorted to prevent such deviations while also

³⁰This interpretation was not possible in the simple one-period model in Section 3, where there was no uncertainty.

³¹Recall that, although 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 .

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 benefit disproportionately high productivity firms. It is then better to not incentivize R&D investments as much for the lower productivity firms, since this makes their allocations more attractive to high productivity firms. In sum, since in this case 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 better.

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, on 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 monopoly valuation term would be increasing in firm type as well.

In addition, as 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.

5.1.2 Age profile of optimal policies

The optimal policies will generically change over firms' life cycles. Age patterns can in principle come from three sources, which are conceptually very different.

The first reason for time dependent policies would be if firms have a finite lifecycle, i.e., if the max age is $T < \infty$. This leads to life cycle considerations such as the shorter horizon for any investments made later in firms' lives. Here the relevant issue is the distance of the period under consideration to the final period T . Both the laissez-faire and the socially optimal investments would naturally decline over a firm's life-cycle, all else equal, as earlier investments contribute to research productivity for more periods. If the technology spillover is positive, as seems natural, the Pigouvian correction term is always positive and, all else constant, will decline over time as the horizon shortens. This age-driven channel is fully eliminated by letting the horizon go to infinity, as we do in our benchmark case. In Online Appendix OA.2, we provide the optimal policies with a finite life cycle.

The second, more fundamental, reason is that the state-contingent policy is set at age 1 under full commitment from the planner. As a result, it is the distance 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 for instance with a first-order autoregressive or geometric autoregressive process with persistence parameter $\tilde{p} < 1$).

This decay towards zero is faster the lower the persistence in types. From the perspective of period 1, as 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.³²

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 as explained 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 $\rho_{\theta r} < \rho_{lr}$.

Finally, there may be direct age effects if the technological fundamentals, such as the step size λ_t , the cost of effort $\phi_t(l)$, and the cost of R&D $M_t(r_t)$ vary with age. One may imagine that as a firm gains expertise, the cost of unobservable and observable R&D inputs may decrease (hence, ϕ_t and M_t would be decreasing with t). More empirical work could shed light on the lifecycle 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. With an infinite lifecycle and age-independent technological parameters, only the second channel applies and wedges converge over time to the Pigouvian correction and monopoly valuation terms. The quantitative analysis in Section 6 will shed light on these age patterns.

5.2 Extensions

Heterogeneity in Production Efficiency: Suppose that firms are also heterogeneous in their production productivities, denoted by θ^p , with realization θ_t^p and history $\theta^{p,t}$. For instance, production costs could be $C(k, \bar{q}_t, \theta_t^p)$. 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 $\tilde{Y}(q_t(\theta^t, \theta^{p,t}), \bar{q}_t, \theta_t^p)$ and profits are $\pi(q_t(\theta^t, \theta^{p,t}), \bar{q}_t, \theta_t^p)$. Similarly to before, we can define $\Pi_t(\theta^t, \theta^{p,t}) := \left(\sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \pi(q_s(\theta^s, \theta^{p,s}), \bar{q}_s, \theta_s^p)}{\partial q_s} \right)$ as the marginal impact on future profit flows from an increase in quality. Let $Q_t(\theta^t, \theta^{p,t}) = \sum_{s=t}^{\infty} \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \tilde{Y}(q_s(\theta^s, \theta^{p,s}), \bar{q}_s, \theta_s^p)}{\partial q_s}$ be the marginal impact of quality on on future expected output net of production costs.

³²In Appendix OA.3, we consider many different cases with regard to the persistence or stochastic process for productivity.

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 firms 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 \(1998\)](#) show 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.

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:

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

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 (25) 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, $\rho_{lr,t}^j - \rho_{\theta r,t}^j$.

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.

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 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 et al. (2016), 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 (25) 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).

5.3 Implementation through Taxes and Subsidies

In this section, we show that the optimal allocations can be implemented with a relatively parsimonious tax function.³³ 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 a simpler implementation of the optimal mechanism, which does not need to 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 iid shocks Albanesi and Sleet (2006).

Market Structure. The constrained efficient allocations solved for in Section 4 are independent of the underlying market structure as long as the information set and toolbox of the planner is as specified there.³⁴ 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 in place, since the latter determines the level of profits in the 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 , i.e., $T_t(q_t, r_t, q_{t-1}, r_{t-1}, q_1)$. The proof is in

³³Until 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 would now like to step away from reporting of 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, i.e., 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.

³⁴For 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, e.g., if firms could save in a hidden way, then the constrained efficient allocation would be different.

the Appendix. Note that because profits are an immediate function of quality q_t , the tax function could instead be written as a function of profits, i.e.,: $\tilde{T}_t(\pi_t, r_t, \pi_{t-1}, r_{t-1}, \pi_1)$.

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 ones. However, this intuition is not correct. Even with Markov shocks, most dynamic tax problems (Farhi and Werning, 2013; Stantcheva, 2017) require that the full past histories are conditioned on. What is different here is that the past stock of quality can serve as a sufficient statistic for the past investments, in the sense that it fully determines the future benefit from more innovation investments (together with r_t and r_{t-1}).

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 .

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

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

Generically, the optimal wedges in Proposition 2 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 non-linearity of the profit tax and R&D subsidy depend 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 A.2, we work out a very simple functional form example that generates constant marginal profit taxes and constant R&D subsidies, i.e., implementation policies that are linear and separable. The latter only depend on the degree of market power (β) and the strength of the spillover (ζ).³⁵

5.4 Controlling Quantity

If the government were also able to intervene in the output market and control the quantity produced, the planning problem is identical to P in (21), except that the impact of quality improvements on the net output produced by the monopolist, $\tilde{Y}(q_t(\theta^s), \bar{q}_s)$, is replaced everywhere

³⁵Estimating the model in the real data, we will show quantitatively below that separable, nonlinear profit taxes and R&D subsidies is sufficient to capture most of the gains from the full mechanism. In fact, a linear profit tax coupled with a nonlinear R&D subsidy does almost equally well.

with the impact of quality improvement on net output as would optimally be chosen by the planner, for every quality level, i.e. $\tilde{Y}^*(q_t(\theta^s), \bar{q}_s)$. Accordingly, in the optimal wedge formulas in Proposition (2), $Q_{t+1}(\theta^{t+1})$ is replaced by $Q_{t+1}^*(\theta^{t+1})$. All else equal, when the planner is also able to control quantity, wedges are larger because the planner is able to make the firm deviate more relative to what it would (suboptimally) do in the laissez-faire. Because being able to control quantity implies removing a constraint in the planner's problem, total output net of all costs will be higher than when quantity cannot be controlled.

6 Quantitative Investigation

In this section and next, we provide empirical content to the theoretical model, by estimating it and numerically illustrating the optimal policies. We first list the open issues for the quantitative analysis relative to the theory:

Open Questions for the Quantitative Analysis

1. What are the estimated values of the key parameters such as the complementarity between R&D and firm research productivity $\rho_{\theta r}$, the persistence \tilde{p} , and the externality strength ζ in the data?
2. What are the cross-sectional patterns of the optimal policies and allocations?
3. What are their age-patterns?
4. By how much do the optimal policies and allocations change when quantity cannot be controlled?
5. How strongly do these key parameters $\rho_{\theta r}$, p , and ζ affect optimal policies?
6. What are the quantitative losses from simpler policies relative to the fully optimal mechanism?

6.1 Data and Summary Statistics

The theory developed in this paper can be applied to different datasets and the model could of course be estimated for different countries, industries, or types of firms to yield different quantitative answers to questions (1)-(6) listed right above and 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 et al. \(2001\)](#)), containing over three million patents with their forward citations. This data match is done and used by [Acemoglu et al. \(2018\)](#) and [Akcigit and Kerr \(2018\)](#). We take several existing data moments as our targets directly from those papers and compute the new ones

directly. To build confidence in our model and estimation, we also match several non-targeted moments in Table 4.

Our code is very flexible and allows for different functional forms and parameter values, as illustrated in the many robustness checks and sensitivity analyses that we describe in Section 6.5. Given the difficulty in solving dynamic mechanism design models with spillovers, we hope that future researchers can use and build on our code.

An alternative dataset would be firm-level accounting data from COMPUSTAT data matched to the NBER patent database. As argued in Bloom, Schankerman, and Van Reenen (2013), and one of the reasons, this data has been widely used, is that these firms represent a large fraction of the innovation in the U.S. For completeness, we provide the numerical results based on the COMPUSTAT sample in the Online Appendix OA.1.

Map between the model and the data: The great advantage of the patent data matched to firm-level data is that there is a natural map 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 , i.e., 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 $q_t = (1 - \delta)q_{t-1} + \lambda_t$. Profits and sales can also directly be measured.

6.2 Estimation

To estimate the model presented in Section 2, we parameterize it as summarized in Table 1. Some of the parameters are calibrated exogenously, following the earlier innovation literature. This reduces the size of the parameter vector to be estimated. These parameters are reported in the upper panel of Table 2. In Section 6.5 we provide sensitivity analyses on each of them. The lower panel of that table reports the key and specific parameters of our model, which are estimated to best match important moments in the data presented in Table 3. In Table 4, we check that the estimation matches well 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 U.S. We call *status quo* economy the economy with the primitives just presented, but in which policies are not optimally set, but rather set to mimic their levels 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 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. The details for this computation are in the Online Appendix. We similarly estimate the effective average corporate tax rate. We find an average

effective subsidy rate of 19% and an average effective corporate tax rate of 23%. Reassuringly, the estimation of the parameters is not very sensitive to the choice of these effective rates within reasonable ranges.

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, $\rho_{\theta l} = \rho_{lr} = 1$. This functional form is specified for tractability; given the data, it would be very difficult to also discipline empirically ρ_{lr} . Given that the sign of $\rho_{lr} - \rho_{\theta r}$ determines the sign of the screening term in the optimal R&D subsidy (as shown in Proposition 2), the key question for whether screening will lead to a higher or lower subsidy on R&D will be whether in the data $\rho_{\theta r} \geq 1$ or $\rho_{\theta r} < 1$. 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} . Appendix Figure OA13 shows the results for a different productivity process, as described in Section 6.5. The shock ε_t follows a normal distribution with mean zero and variance σ_ε .

TABLE 1: 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	$C_t(k, \bar{q}_t)$	$\frac{k}{\bar{q}_t^\zeta}$
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	$\phi_t(l_t)$	$\kappa_l \frac{l_t^{1+\gamma}}{1+\gamma}$
Cost of R&D	$M_t(r_t)$	$\kappa_r \frac{r_t^{1+\eta}}{1+\eta}$
Stochastic type process	$f^t(\theta_t \theta_{t-1})$	$\log \theta_t = \tilde{p} \log \theta_{t-1} + (1 - \tilde{p})\mu_\theta + \varepsilon_t$
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

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

Externally Calibrated Parameters: We take the externally calibrated parameters from reputable papers in the innovation and growth literatures, but we are also careful to provide a lot of sensitivity analysis on their values in Section 6.5.

The profit parameter β is set to 0.15 as in Guner et al. (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 described in detail in Hall et al. (2005), where the value of 10% is also taken from. 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 from the data. 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 3 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 of the model.

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

$$\chi = (\alpha, \rho_{\theta r}, \sigma_\varepsilon, p, \kappa_l, \kappa_r, \gamma, \zeta, \Theta^1)$$

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

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

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 Tables OA2 and OA3, as well as Figures OA11 and OA12 we provide the results 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 in generating successful innovations. It has been estimated widely 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 is computed by Acemoglu et al. (2018) and is consistent with other papers too. 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 highlighted the role of firm heterogeneity and the role of uncertainty over time.

TABLE 2: 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	$\tilde{\rho}$	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_1	Θ^1	1.91	(0.097)
Production externality	ζ	0.018	(0.001)

TABLE 3: 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)

Moment four hence 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 $\tilde{\rho}$ 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. “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). As is intuitive, 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, moment 8, targets the estimate of technological spillovers in Bloom, Schankerman, and Van Reenen (2013). These authors estimate 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 instrument 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). The fit we obtain is very close. This process helps us identify the externality 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. This has been estimated in many papers and we pick one of the most respected sources, which is Bloom et al. (2002) who find an elasticity of R&D to user costs of -0.35. Our model matches this very well.

Goodness of Fit: Non-targeted moments. To check whether the fit of our estimated model is good even for non-targeted moments, we provide in Table 4 the values of four important and non-targeted moments in the data and the model, which pertain to the lifecycle 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 bottom 90% versus top 10% firms. The fit is quite good, lending further credibility to our estimation and confirming that we are able to also capture the more detailed tail behavior of the data.

TABLE 4: 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

6.3 Results

Table 2 shows the estimated parameters of the model. Focusing on some of the key parameters that were highlighted in Section 5, 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, presented in analytical form in Section 5.

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's illustrate the difference with linear taxes, to make it simpler. If the profit tax applies to profits gross of R&D spending, i.e., 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:

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

With a gross subsidy \tilde{s} , the net incentive is not captured by the subsidy rate itself, since the profit tax captures part of the return to R&D investments. Thus, a share of the gross subsidy simply goes towards cancelling out that 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.

Thus, in the figures below we also show the “gross” R&D wedge $\tilde{s}(\theta^t)$, namely the gap between marginal costs and marginal benefits of R&D, taking into account the R&D effort wedge, i.e., that there is simultaneously a tax on profits, in addition to depicting $s(\theta^t)$:³⁶

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

³⁶ $\tilde{s}(\theta^t)$ is related to $s(\theta^t)$ from Proposition 2 through

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

In the linear example above, s is directly comparable to the average wedge $s(\theta^t)$ from Section 5, while \tilde{s} is comparable to $\tilde{s}(\theta^t)$.

6.3.1 Cross-sectional Patterns of the Optimal Allocations

We now turn to answering question (2). Panels C and D of Figure 2 plot the optimal profit wedge $\tau(\theta^t)$, the gross R&D wedge $\tilde{s}(\theta^t)$, and the corresponding net R&D wedge $s(\theta^t)$ for firms of different profit levels and R&D investments for ages $t = 2, 5, \text{ and } 15$. Panels E and F also represent 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 otherwise would in the *laissez-faire*. Thinking back of the optimal wedge formulas in Section 5, this is to account for the monopoly distortion effect (incentivize monopolists to produce more indirectly) and the Pigouvian correction effect (i.e., 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's consider how wedges vary by firm type, levels of profits, and R&D expenses, remembering the theoretical discussion in Section 5. 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 described before: since it is higher type firms that are tempted to pretend to be lower types, lower types firms' allocations are distorted to prevent this from happening. 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 $\rho_{\theta r} > 1$, and so that 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 benefit disproportionately high productivity firms. It is better to incentivize R&D investments less for the lower productivity firms, as this makes mimicking more attractive for high productivity firms.³⁷ In short, lower productivity firms which have no comparative advantage at innovation here are not incentivized as much to invest in R&D, so that high productivity firms can be incentivized better. 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

³⁷Recall 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*. In the *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.

pattern when plotted against observables, namely profits (panel C) and R&D investments (panel D). It is worth reiterating that what summarizes the incentive provision for innovation is the net wedge; the gross wedge is depicted as it matches better the intuitions that comes 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 nevertheless, 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 them 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 7, 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 a lower marginal profit taxes on more profitable firms and lower marginal subsidy at higher levels of R&D investments. But, quantitatively, other implementations may work similarly well, if the loss from fine-tuning is small. For instance, a constant profit tax that is more generous than it should 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 a too generous tax is quantitatively small.

Panel C and D of Figure 3 shows 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 $\rho_{l\theta} > 0$ and $\rho_{r\theta} > 0$, effort and R&D expenses of 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.

6.3.2 Age Patterns of the Optimal Allocations

As explained in Section 5, age patterns can in principle arise for three reasons: a finite life cycle horizon, the fact that screening policies are set at age one with full commitment, and a possible age-dependency of the primitives of the model. In the current setting, the life cycle is infinite and there are no age-dependent primitives per se, so that the age patterns are purely driven by the screening patterns. Panels A and B from Figure 2 answer question (3) listed at the beginning of this Section by plotting 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 2 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.

Figures OA9 and OA10 show what happens when the life cycle is finite, with a given death and exit rate. In this case, the age paths of optimal inputs are hump-shaped, driven by the balance of the screening considerations and the life cycle considerations. In the first part of the life cycle, the screening considerations dominate; in the latter part, the dominant forces are the finite life cycle and the approach of the terminal period, which make investments less lucrative, privately and socially. Thus, with a finite life cycle, young firms, up to mid-life, should optimally provide an increasing amount of effort and investments for R&D. After mid-life, the effort and investment are declining given the shortening horizon left to reap the benefits.

6.3.3 Optimal Allocations and Wedges When Quantity Can be Controlled

When quantity can also be controlled, the planner has an additional lever that can also be made part of the contract. As a result, the planner can make firms deviate even more from their laissez-faire allocations to induce a better allocation. Accordingly, the wedges are larger in absolute value, as illustrated in Figure 4. Overall, the innovation inputs and step sizes are larger, as shown in Figure 5.

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

In Figure 6, we quantify the effect of the key parameters (question (5)). 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, and the flatter the net wedge is for different ages, as it is then more and more driven by only the monopoly valuation and Pigouvian correction terms.

The persistence of the firm's research productivity process affects the optimal policies very significantly, in particular their age pattern. Panels C and D depict the wedges for a higher value of persistence than our benchmark one, 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. Figure OA13 shows the wedges for a first-order autoregressive process instead of the benchmark logarithmic autoregressive process. Figure OA15 shows the wedges with an increasing persistence over the life cycle that has the same average value as our benchmark estimate ($p = 0.63$). Since persistence is now higher later in life, the decay of the wedges is faster. Figures OA17 and OA19, respectively, show the full set of wedges for a lower persistence ($p = 0.5$) and a higher persistence ($p = 0.9$) than in our estimation. The speed of convergence to the Pigouvian and monopoly correction terms is strongly shaped by the type process.

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 investment wedges.

6.5 Robustness Checks and Sensitivity Analysis

We provide many robustness checks and sensitivity analyses in the Appendix. In section OA.1, we re-estimate the model only on publicly-traded firms from COMPUSTAT matched to patent data. There isn't a very sizable difference in the optimal wedges or allocations because many of the innovating and patenting firms do appear in COMPUSTAT at some point. In section OA.2, we show what happens with a finite firm life cycle, in which case the horizon becomes important.

In Figures OA11 and OA12, we perform a type of two-step GMM estimation with weights taken from the variance-covariance matrix of moments. The reason this is not our benchmark is because we do not have the full variance-covariance matrix as moments M8 and M9 are taken from other papers (based on good identification strategies, e.g. to identify spillovers). We hence assume the off-diagonal terms are zero. Table OA3 shows the match for the targeted moments and Table OA2 the estimated parameter values. The results are very similar to our benchmark ones.

In the remaining figures we change the externally calibrated parameters. In all these cases, it is important bearing in mind that wedges represent the gap between what firms would do in the laissez-faire and what the planner induces them to do in the optimal mechanism. Variations in any of these parameters not only change the optimal allocation, they also change what firms would optimally do in the laissez-faire, often in the same direction. As a result, the wedges may not change that much from a change in these parameters; however, the allocations induced could be very different. This is why we show all the wedges and the allocations for each set of parameter values.³⁸

In Figures OA13 to OA20, we explore the role of the stochastic type process assumed, some of which was already discussed above. More precisely, Figures OA13 and OA14 show the wedges for a first-order autoregressive process; Figure OA15 and OA16 an increasing persistence over

³⁸In addition, total revenues raised by the government and consumer welfare would also be very different since it depends not just on the total innovation produced, but also on the share that goes to consumers.

the life cycle; Figures OA17 to OA20 respectively have $p = 0.5$ and $p = 0.9$. The persistence of this stochastic process affects the rate of decay of the wedges very significantly, but not the qualitative findings described above. In addition, a more persistent process increases the ability of the planner to provide dynamic incentives and improves the allocations: there are higher levels of effort and R&D investment for firms of all productivities.

Figures OA21 to OA24 show the changes induced by higher or lower values of β . Higher β represents a higher degree of market power, as it increases the markup over marginal costs that the intermediate good producer can charge. At the same time, it 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.

Figures OA25 and OA28 consider higher rates of depreciation of innovation, of $\delta = 0.15$ and $\delta = 0.3$ respectively. The higher the rate of depreciation, the higher the wedges have to be to induce firms to invest sufficiently much (relative to what they would do if left to choose). Naturally, the higher the rate at which knowledge depreciates and the lower the optimal investments, step sizes, and resulting innovation that can be stimulated.

Finally, Figures OA31 and OA32 show what happens when the cost of R&D is less convex, i.e., when $\eta = 1$. This barely changes the wedges, as they represent the share of costs that is subsidized. However, as can be expected, the level of R&D effort and incentives that can be incentivized are larger when costs are less convex.

7 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 6. We then compute the welfare loss relative to the welfare obtained in the unrestricted mechanism.³⁹

Table 5 shows the results. In this table, each panel considers a separate class of policies, ranging step-by-step from linear to nonlinear and non-separable ones. The three columns show, respectively, the welfare achieved from the optimal policy in each class relative to (i) our benchmark optimal planner solution, (ii) the planning problem in which quantity can be controlled as well, and (iii) 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 U.S., i.e., 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 and 7% relative to the case where quantity can also be controlled. If there were no spillovers at all, current

³⁹Bear in mind that “welfare” here is net output or consumption going to the consumer and has no redistributive component.

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 one and yields 89% of the welfare gain from the full optimum, 92.4% of the optimum with quantity control 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) improves the welfare gains further.

The biggest gain comes from a nonlinear Heathcote-Storesletten-Violante (HSV) policy, as used by Heathcote et al. (2017) to study optimal progressivity and subsequently by Heathcote et al. (2019) to study optimal age-dependent progressivity (panel D). In these papers, 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 it here 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 quantitatively 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. It is clear that the most important gain comes from a nonlinear R&D subsidy: A linear subsidy plus HSV tax system yields only 94.7% of the full gain. On the contrary, linearizing the profit tax generates only a very small welfare loss of 0.1% relative to the fully nonlinear HSV policy. 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 too generous incentives to the low profitability firms (and just about the right level for high profitability firms), that loss is quantitatively very small as they make low profits to start with. This is a particularly useful finding because corporate taxes are typically more or less linear around the world. 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 ones. The reasons to tax or subsidize non-innovating firms would be different and could be nested in our model if spillovers are shut off and the market between intermediate and final good producers

TABLE 5: WELFARE FROM OPTIMAL SIMPLER POLICIES

Policy Type	Welfare Achieved Relative to Full Optimum		
	Benchmark	Quantity control	No spillovers
<i>A. Current Linear</i>			
$T'(\pi) = 0.23$ $S'(M) = 0.19$	18%	7%	31.1%
<i>B. Optimal Linear</i>			
$T'(\pi) = \tau_0$ $S'(M) = s_0$	89%	92.4%	88.5%
<i>C. Linear with Interaction Term</i>			
$T'(\pi, M) = \tau_0 + \tau_1 M$ $S'(M) = s_0$	93.5%	95.1%	93.7%
<i>D. Heathcote-Storesletten-Violante (HSV)</i>			
$T'(\pi) = \tau_0 - \tau_1 \pi^{\tau_2}$ $S'(M) = s_0 - s_1 M^{s_2}$	97.4%	96.3%	98.2%
<i>E. HSV tax on profits and linear subsidy</i>			
$T'(\pi) = \tau_0$ $S'(M) = s_0 - s_1 M^{s_2}$	94.7%	95.8%	95.6%
<i>F. HSV subsidy on R&D and linear profit tax</i>			
$T'(\pi) = \tau_0$ $S'(M) = s_0 - s_1 M^{s_2}$	97.3%	96.2%	97.4%
<i>G. HSV with Interaction Term</i>			
$T'(\pi, M) = \tau_0 + \tau_3 M^{s_2} - \tau_1 \pi^{\tau_2}$ $S'(M) = s_0 - s_1 M^{s_2}$	97.4%	96.4 %	98.3 %

Notes: 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 (2) relative to the optimum when quantity can be controlled; Column (3) for the benchmark optimum but when there is no spillover $\zeta = 0$).

is calibrated to match the one for non-innovating firms. Our framework for firm taxation is malleable and quite general and we think 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 at all 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.

8 Conclusion

In this project, 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. Firms' efficiency in converting research inputs into research outputs, as well as an important input into the innovation process (called "R&D effort") are unobservable to the government. Policies should ideally be targeted towards the most efficient firms, but asymmetric information makes this more challenging. We overcome this challenge by using and expanding a combination of mechanism design techniques developed in the recent contract theory and new dynamic public finance literatures that we augment with spillovers. We combine elements from three literatures: a macro innovation model, the theoretical tools of mechanism design, and an estimation using micro-level firm data matched to patent data to discipline and quantify our model. This type of dynamic asymmetric information model with spillovers and the solution method could be applied more broadly to other types of firm investments with or without externalities, as well as to firm taxation more generally.

We characterize the constrained efficient allocations that arise in a direct revelation mechanism with spillovers, which does not impose any *ex ante* restriction on the policy tools the government can use. We find that 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. We highlight that a crucial statistic is the complementarity of R&D investments to R&D effort (i.e., the complementarity between observable and unobservable innovation inputs) relative to the complementarity of R&D investments to unobservable firm research productivity: the more complementary R&D investment is to firm research productivity, the more rents a firm can extract if R&D investments are subsidized. Screening considerations can hence dampen the first-best corrective policies. The persistence of firm 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.

The estimation of our model based on key moments in the data allows us to quantify the op-

timal 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 any given R&D investment. Since higher productivity firms have a comparative advantage at innovation, it is better to incentivize R&D investments less for the lower productivity firms, as this makes mimicking them more attractive for high productivity firms. We discuss how these incentives translate into wedges, first, and then into actual taxes and subsidies.

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. The intuition for why this is the best screening mechanism is that 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 going to be more attractive to them than it will be to low productivity firms.

Regarding the 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 and thus perfectly mimics the shape of the wedges. In fact, quantitatively, the most important feature is the nonlinearity in the R&D subsidy; making the profit tax 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 giving low profit firms a too generous tax is quantitatively small given that they make low profits to start with. Thus, linear corporate income taxes— as they are often seen around the world — can be very close to optimal for innovating firms if combined with the right nonlinear R&D subsidy.

We hope that future research could 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 the following and many other 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, it would be very interesting to study optimal R&D policies when there is a noisy signal about product quality which may be manipulable by firms. 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.

Our framework provides a new way of thinking about both innovation policy and firm taxation more generally. Introducing asymmetric information in the heterogeneous and stochastic firm types allows to capture many features of the real world. The framework can be extended in many directions to address important questions in policy design for firms.

References

- Aboody, D. and B. Lev (2000). Information Asymmetry, R&D, and Insider Gains. *The Journal of Finance* 55(6), 2747–2766.
- Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018). Innovation, reallocation and growth. *American Economic Review* 108(11), 3450–3491.
- Aghion, P. and P. Howitt (1992). A model of growth through creative destruction. *Econometrica* 60(2), 323–351.
- Akcigit, U., J. Grigsby, T. Nicholas, and S. Stantcheva (2018). Taxation and innovation in the 20th century. Working Paper 24982, National Bureau of Economic Research.
- Akcigit, U., D. Hanley, and N. Serrano-Velarde (2016). Back to Basics: Basic Research Spillovers, Innovation Policy and Growth. CEPR Discussion Paper 11707.
- Akcigit, U. and W. R. Kerr (2018). Growth through Heterogeneous Innovations. *Journal of Political Economy* 126(4), 1374–1443.
- Alam, P. and K. S. Walton (1995). Information Asymmetry and Valuation Effects of Debt Financing. *The Financial Review* 30(2), 289–311.
- Albanesi, S. and C. Sleet (2006). Dynamic Optimal Taxation with Private Information. *Review of Economic Studies* 73(1), 1–30.
- Ales, L., S.-H. Cho, and E. Korpeoglu (2017). Optimal Award Scheme in Innovation Tournaments. *Operations Research* 65(3), 693–702.
- Ales, L., M. Kurnaz, and C. Sleet (2015). Technical Change, Wage Inequality and Taxes. *American Economic Review* 105(10), 3061–3101.
- Ales, L. and C. Sleet (2016). Taxing Top CEO Incomes. *American Economic Review* 106(11), 3311–3366.
- Baron, D. P. and R. B. Myerson (1982). Regulating a Monopolist with Unknown Costs. *Econometrica* 50(4), 911–930.
- 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. In *The Regulation of Science and Technology*, pp. 136–160. Springer.
- Bloom, N. and R. Griffith (2001). The Internationalisation of UK R&D. *Fiscal Studies* 22(3), 337–355.
- Bloom, N., R. Griffith, and J. Van Reenen (2002, July). Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979–1997. *Journal of Public Economics* 85(1), 1–31.
- Bloom, N., E. A. Mahajan, D. McKenzie, and J. Roberts (2013). Does Management Matter? Evidence from India. *Quarterly Journal of Economics* 128(1), 1–51.
- Bloom, N., R. Sadun, and J. V. Reenen (2012). The Organization of Firms across Countries. *Quarterly Journal of Economics*.
- Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying Technology Spillovers and Product Market Rivalry. *Econometrica* 81(4), 1347–1393.
- Bloom, N. and J. Van Reenen (2007). Measuring and Explaining Management Practices Across Firms and Countries. *Quarterly Journal of Economics*.

- Bloom, N., J. Van Reenen, and H. Williams (2019). A Toolkit of Policies to Promote Innovation. *Journal of Economic Perspectives* 33(3), 163–184.
- Chari, V., M. Golosov, and A. Tsyvinski (2012). Prizes and Patents: Using Market Signals to Provide Incentives for Innovations. *Journal of Economic Theory* 147, 781–801.
- Chen, Z., Z. Liu, J. Suárez Serrato, and D. Xu (2019). Notching R&D Investment with Corporate Income Tax Cuts in China. *Working paper (under review)*.
- Cullen, J. and R. Gordon (2006). Tax Reform and Entrepreneurial Activity. *Tax Policy and the Economy* 20, 41–72.
- Cullen, J. and R. Gordon (2007). Taxes and Entrepreneurial Risk Taking: Theory and Evidence for the U.S. *Journal of Public Economics* 91(7-8), 1479–1505.
- Farhi, E. and I. Werning (2013). Insurance and Taxation over the Life Cycle. *Review of Economic Studies* 80(2), 596–635.
- Gallini, N. and S. Scotchmer (2002). Intellectual Property: When is it the Best Incentive Mechanism? *Innovation Policy and the Economy* 2, 51–78.
- 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.
- Golosov, M., A. Tsyvinski, and I. Werning (2006). New Dynamic Public Finance: A User's Guide. In *NBER Macroeconomics Annual 2006, Volume 21*, pp. 317–388. MIT Press.
- Golosov, M., A. Tsyvinski, and N. Werquin (2014). A variational approach to the analysis of tax systems. Working Paper 20780, NBER.
- Gompers, P. A. (1995). Optimal Investment, Monitoring, and the Staging of Venture Capital. *The Journal of Finance* 50(5), 1461–1489.
- Goolsbee, A. (1998). Does r&d policy primarily benefit scientists and engineers? *American Economic Review (Papers and Proceedings)* 88(2), 298–302.
- Gordon, R. and Y. Lee (2005). Tax Structure and Economic Growth. *Journal of Public Economics* 89(5-6), 1027–1043.
- Gordon, R. and Sarada (2018). How Should Taxes be Designed to Encourage Entrepreneurship? *Journal of Public Economics* 166, 1–11.
- Green, J. and S. Scotchmer (1995). On the Division of Profit in Sequential Innovation. *Journal of Industrial Economics* 26, 20–33.
- Griliches, Z. (1998). *R&D and Productivity: The Econometric Evidence*. University of Chicago Press.
- Grossman, V., T. Steger, and T. Trimborn (2013). Dynamically optimal r&d subsidization. *Journal of Economic Dynamics & Control* 37, 516–534.
- Guner, N., G. Ventura, and Y. Xu (2008). Macroeconomic Implications of Size-Dependent Policies. *Review of Economic Dynamics* 11, 721–744.
- Hall, B., A. Jaffe, and M. Trajtenberg (2001). The nber patent citations data file: Lessons, insights and methodological tools. Working Paper 8498, National Bureau of Economic Research.
- Hall, B., A. Jaffe, and M. Trajtenberg (2005). Market Value and Patent Citations: A First Look. *RAND Journal of Economics* 36(1638).
- 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.

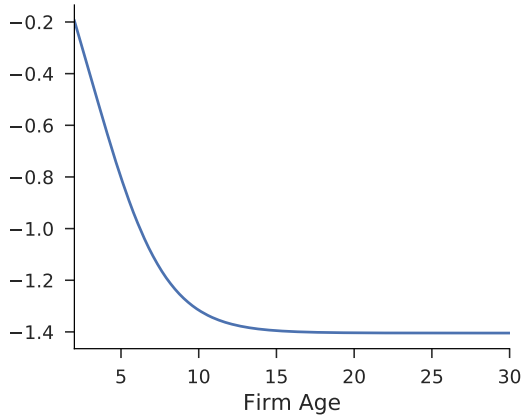
- Hall, B. H. and J. Lerner (2009, 09). The financing of r&d and innovation. *NBER Working Paper Series*, 15325. Copyright - Copyright National Bureau of Economic Research, Inc. Sep 2009; Last updated - 2015-12-18.
- Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal Tax Progressivity: An Analytical Framework. *The Quarterly Journal of Economics* 132(4), 1693–1754.
- Heathcote, J., K. Storesletten, and G. L. Violante (2019). Optimal progressivity with age-dependent taxation. Working Paper 25617, National Bureau of Economic Research.
- Hicks, J. (1970). Elasticity of Substitution Again: Substitutes and Complements. *Oxford Economic Papers* 25, 289–296.
- Hopenhayn, H., G. Llobet, and M. Mitchell (2006). Rewarding sequential innovators: Prizes, patents, and buyouts. *Journal of Political Economy* 114(6), 1041–1068.
- Hopenhayn, H. and M. Mitchell (2001). Innovation, Variety and Patent Breadth. *The RAND Journal of Economics* 32, 152–166.
- Hopenhayn, H. and M. Mitchell (2011). Optimal Patent Policy with Recurrent Innovators. *UCLA Working Paper*.
- Hsieh, C.-T. and P. Klenow (2009). Misallocation and manufacturing tfp in china and india. *Quarterly Journal of Economics* 124, 1403–1448.
- Hsieh, C.-T. and P. Klenow (2014). The life cycle of plants in india and mexico. *Quarterly Journal of Economics* 129, 1035–1084.
- Jones, C. I. and J. C. Williams (1998). Measuring the social return to r&d. *Quarterly Journal of Economics* 113(4), 1119–1135.
- Jones, C. I. and J. C. Williams (2000). Too much of a good thing? the economics of investment in r&d. *Journal of Economic Growth* 5, 65–85.
- Kaplan, S. N. and P. Stromberg (2001, 05). Venture capitalists as principals: Contracting, screening, and monitoring. *The American Economic Review* 91(2), 426–430. Copyright - Copyright American Economic Association May 2001; Last updated - 2016-04-30; CODEN - AENRAA; SubjectsTermNotLitGenreText - United States; US.
- Kerr, W. R., R. Nanda, and M. Rhodes-Kropf (2014). Entrepreneurship as experimentation. *Journal of Economic Perspectives* 28(3), 25–48.
- Kindermann, F. and D. Krueger (2016). High marginal tax rates on the top 1%? lessons from a life cycle model with idiosyncratic income risk. Working paper, University of Pennsylvania.
- Klemperer, P. (1990). How Broad Should the Scope of Patent Protection Be? *Journal of Industrial Economics* 21, 113–130.
- Kleven, H. (2004). Optimum Taxation and the Allocation of Time. *Journal of Public Economics* 88, 545–557.
- Kleven, H., C. Kreiner, and E. Saez (2009). The Optimal Income Taxation of Couples. *Econometrica* 77, 537–560.
- Laffont, J.-J. and J. Tirole (1986). Using Cost Observation to Regulate Firms. *Journal of Political Economy* 94(3), 614–641.
- Lockwood, B. and M. Weinzierl (2015). De gustibus non est taxandum: Theory and evidence on preference heterogeneity and redistribution. *Journal of Public Economics* 124, 74–80.

- Lockwood, B. B. (2017). Optimal income taxation with present bias. Working paper, Wharton.
- Lockwood, B. B., C. G. Nathanson, and E. G. Weyl (2017). Taxation and the Allocation of Talent. *Journal of Political Economy* 125(5), 1635–1682.
- Mansfield, E. (1986). The R&D Tax Credit and Other Technology Policy Issues. *The American Economic Review* 76(2), 190–194.
- Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica* 70(2), 583–601.
- O'Donoghue, T. and J. Zweimuller (2004). Patents in an Model of Endogenous Growth. *Journal of Economic Growth* 9, 81–123.
- Pakes, A. (1985, April). On Patents, R&D, and the Stock Market Rate of Return. *Journal of Political Economy* 93(2), 390–409.
- Pakes, A. (1986). Patents as Options: Some Estimates of the Value of Holding European Patent Stocks. *Econometrica* 54(4), 755–784.
- Pakes, A. and Z. Griliches (1984). Patents and R and D at the Firm Level: A First Look. In Z. Griliches (Ed.), *R and D Patents and Productivity*, pp. 55–72. University of Chicago Press.
- Pavan, A., I. Segal, and J. Toikka (2014). Dynamic Mechanism Design: A Myersonian Approach. *Econometrica* 82(2), 601–653.
- 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), 230–71.
- Rochet, J.-C. and P. Choné (1998). Ironing, Sweeping, and Multidimensional Screening. *Econometrica* 66(4), 783–826.
- Romer, P. (1990). Endogenous technological change. *Journal of Political Economy* 98(5), 71–102.
- Rothschild, C. and F. Scheuer (2013). Redistributive Taxation in the Roy Model. *The Quarterly Journal of Economics* 128(2), 623–668.
- Rothschild, C. and F. Scheuer (2016). Optimal Taxation with Rent-Seeking. *The Review of Economic Studies* 83(3), 1225–1262.
- Sachs, D., A. Tsyvinski, and N. Werquin (2016). Nonlinear Tax Incidence and Optimal Taxation in General Equilibrium. *Working Paper, Toulouse School of Economics*.
- Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies* 68(1), 205–229.
- Sappington, D. (1982). Optimal Regulation of Research and Development under Imperfect Information. *Bell Journal of Economics* 13(2), 354–368.
- Scotchmer, S. (1999). On the optimality of the patent renewal system. *RAND Journal of Economics* 30(2), 181–196.
- Serrano-Velarde, N. (2009). Crowding-Out At The Top: The Heterogeneous Impact of R&D Subsidies on Firm Investment . Bocconi Working Paper.
- Stantcheva, S. (2015). Learning and (or) doing: Human capital investments and optimal taxation. Working Paper 21381, National Bureau of Economic Research.
- Stantcheva, S. (2017). Optimal taxation and human capital policies over the life cycle. *Journal of Political Economy*, 125(6), 1931–1990.

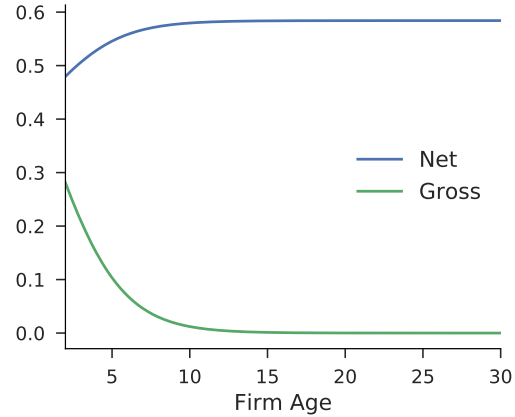
- Syverson, C. (2011). What determines productivity? *Journal of Economic Literature* 49(2), 326?365.
- Terwiesch, C. and Y. Xu (2008). Innovation contests, open innovation and multiagent problem solvings. *Management Science* 54(9), 1529– 1543.
- Tyson, L. and G. Linden (2012). The Corporate R&D Tax Credit and U.S. Innovation and Competitiveness: Gauging the Economic and Fiscal Effectiveness of the Credit. *Washington: Center for American Progress*.
- Weinzierl, M. (2014). The promise of positive optimal taxation: Normative diversity and a role for equal sacrifice. *Journal of Public Economics* 118, 128–142.
- Werquin, N. (2016). Income Taxation with Frictional Labor Supply. Working Paper, Toulouse School of Economics.
- 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.

FIGURE 2: OPTIMAL PROFIT AND R&D WEDGES

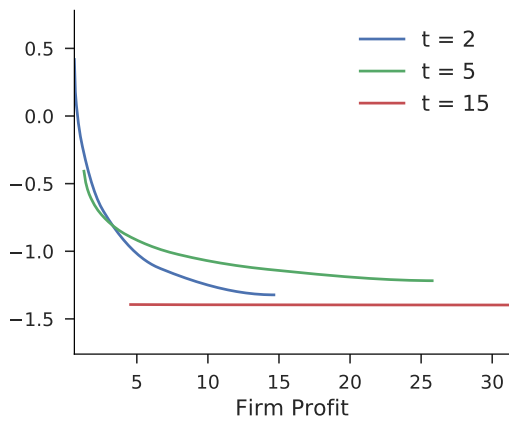
(a) Profit wedge by age



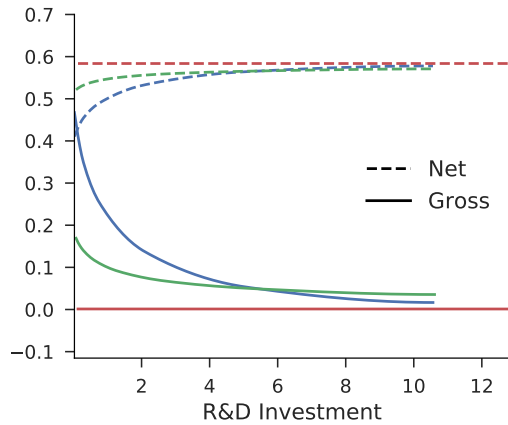
(b) R&D wedges by age




(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



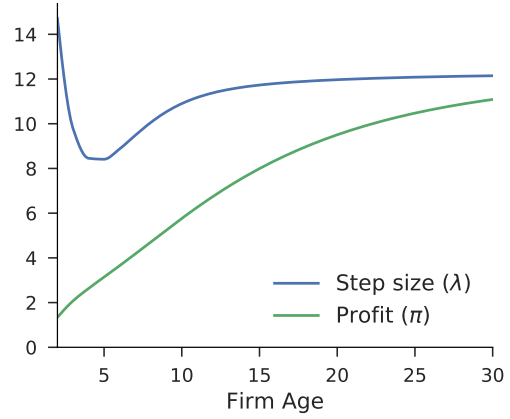
Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE 3: OPTIMAL ALLOCATIONS


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type



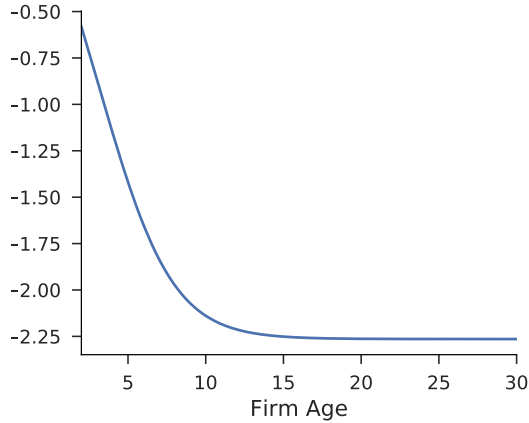
(d) R&D investments by Type



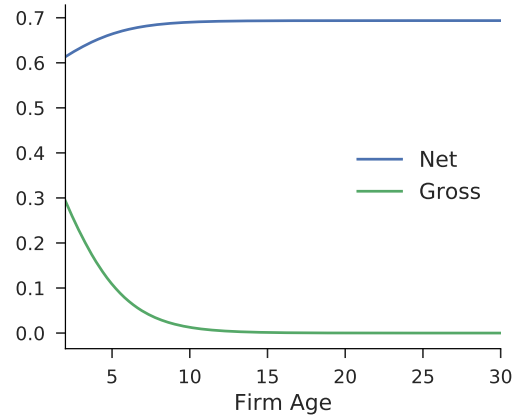
Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE 4: OPTIMAL PROFIT AND R&D WEDGES WITH QUANTITY CONTROL


(a) Profit wedge by age




(b) R&D wedges by age




(c) Profit wedge as function of profits



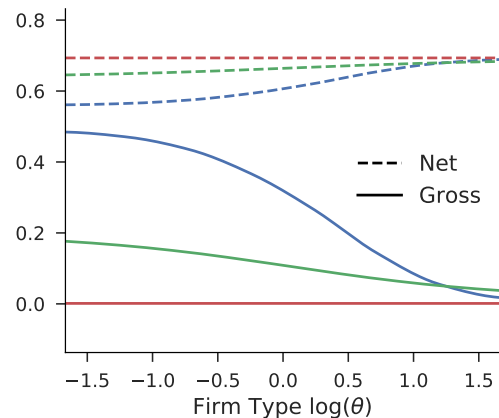
(d) R&D wedges as functions of R&D investments



(e) Profit wedge as function of type θ_t




(f) R&D wedges as functions of type θ_t




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE 5: OPTIMAL ALLOCATIONS WITH QUANTITY CONTROL


(a) Investments and Effort by Age



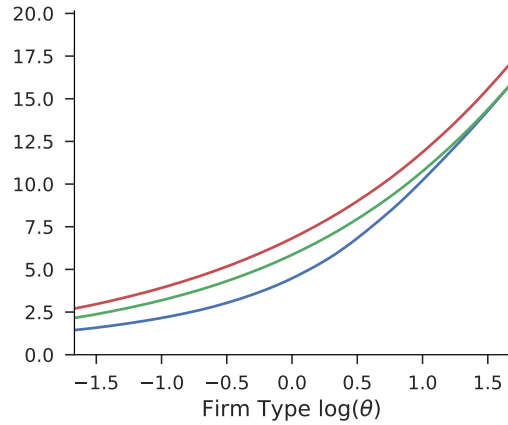
(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type



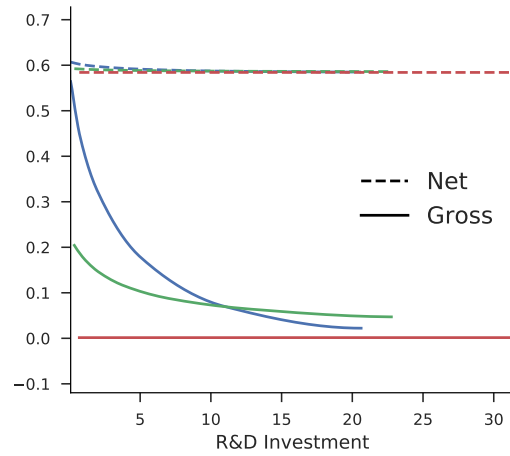
Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE 6: COMPARATIVE STATICS: OPTIMAL PROFIT AND R&D WEDGES


(a) Profit wedge with $\rho_{\theta r} = 0.8$



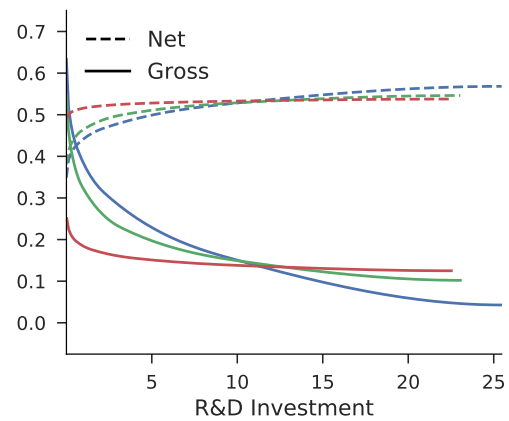
(b) R&D wedges with $\rho_{\theta r} = 0.8$




(c) Profit wedge with $\tilde{p} = 0.9$




(d) R&D wedges with $\tilde{p} = 0.9$



(e) Profit wedge with $\zeta = 0$



(f) R&D wedges with $\zeta = 0$



Notes: Panels (a) and (b) show the wedges for $\rho_{\theta r} = 0.8$. Panels (c) and (d) show the wedges for $\tilde{p} = 0.9$. Panels (e) and (f) show the wedges for $\zeta = 0$.

Appendix

A.1 Additional Results

A.1 Simple Model

We illustrate here the two implementations in the case in which quantity can also be controlled. The benchmark case where quantity cannot be controlled is treated in detail in Section 5.3.

Tax implementation:

First, the government can subsidize the price of production at a nonlinear rate $s_p(k, q)$ as a function of the quantity and quality of the good sold to the final good producer, such that the post subsidy price is $(1 + s_p(k, q))p(k, q) = \frac{Y(k, q)}{k}$, and in addition levy a profit tax (which could be negative) $T(\pi, r)$ that depends nonlinearly on profits and R&D investments. Firms choose quantity to maximize profits conditional on quality, which, thanks to the price subsidy, becomes equivalent to maximizing household consumption net of production costs. Note that under a constant monopoly price markup (as arises for instance under the functional form assumptions in Section 6 where $Y(q, k) = \frac{1}{1-\beta}q^\beta k^{1-\beta}$), the price subsidy needed to align the monopolist's post-tax price with social marginal valuation of quantity is constant and equal to $\frac{\beta}{1-\beta}$. With this price subsidy, profits will be equal to $\tilde{Y}^*(q_0 + \lambda(r, l, \theta_i), \bar{q})$. The maximization problem of a firm of type θ_i with respect to the remaining choices of l and r is then:

$$\max_{l, r} \{ \tilde{Y}^*(q_0 + \lambda(r, l, \theta_i), \bar{q}) - T(\tilde{Y}^*(q_0 + \lambda(r, l, \theta_i), \bar{q}), r) - \phi(l) - M(r) \}$$

The first-order conditions of the firm with this tax implementation are:

$$\begin{aligned} -\frac{\partial T(\tilde{Y}^*(q(\theta_i), \bar{q}), r(\theta_i))}{\partial r(\theta_i)} + \frac{\partial \tilde{Y}^*(q(\theta_i), \bar{q})}{\partial q} \frac{\partial \lambda(r(\theta_i), l(\theta_i), \theta_i)}{\partial r(\theta_i)} \left(1 - \frac{\partial T(\tilde{Y}^*(q(\theta_i), \bar{q}), r(\theta_i))}{\partial \pi} \right) &= M'(r(\theta_i)) \\ \left(1 - \frac{\partial T(\tilde{Y}^*(q(\theta_i), \bar{q}), r(\theta_i))}{\partial \pi} \right) \frac{\partial \tilde{Y}^*(q(\theta_i), \bar{q}), r(\theta_i)}{\partial q} \frac{\partial \lambda(r(\theta_i), l(\theta_i), \theta_i)}{\partial l(\theta_i)} &= \phi'(l(\theta_i)) \end{aligned}$$

We can use the first-order conditions of the firms into the optimal wedge formulas to obtain a characterization of the optimal (explicit) marginal tax and subsidy:

$$\begin{aligned} -\frac{1}{\frac{\partial w(r(\theta_1), \theta_1)}{\partial r} l(\theta_1)} \frac{\partial T(\tilde{Y}^*(q(\theta_1), \bar{q}), r(\theta_1))}{\partial r(\theta_1)} &= \frac{\partial T(\tilde{Y}^*(q(\theta_i), \bar{q}), r(\theta_i))}{\partial \pi} \frac{\partial \tilde{Y}^*(q(\theta_i), \bar{q})}{\partial q} \\ + \left(f_1 \frac{\partial \tilde{Y}^*(q(\theta_1), \bar{q})}{\partial \bar{q}} + f_2 \frac{\partial \tilde{Y}^*(q(\theta_2), \bar{q})}{\partial \bar{q}} \right) + \frac{f_2}{f_1} \left(1 - \frac{\frac{\partial \log(w(r(\theta_1), \theta_2))}{\partial \log(r)}}{\frac{\partial \log(w(r(\theta_1), \theta_1))}{\partial \log(r)}} \right) &\frac{1}{w(r(\theta_1), \theta_2)} \phi' \left(\frac{w(r(\theta_1), \theta_1) l(\theta_1)}{w(r(\theta_1), \theta_2)} \right) \end{aligned}$$

$$\begin{aligned} \frac{\partial T(\tilde{Y}^*(q(\theta_i), \bar{q}), r(\theta_i))}{\partial \pi} \frac{\partial \tilde{Y}^*(q(\theta_i), \bar{q}), r(\theta_i))}{\partial q} &= - \left(f_1 \frac{\partial \tilde{Y}^*(q(\theta_1), \bar{q})}{\partial \bar{q}} + f_2 \frac{\partial \tilde{Y}^*(q(\theta_2), \bar{q})}{\partial \bar{q}} \right) \\ &\quad - \frac{f_2}{f_1} \left(\frac{1}{w(r(\theta_1), \theta_2)} \phi' \left(\frac{w(r(\theta_1), \theta_1) l(\theta_1)}{w(r(\theta_1), \theta_2)} \right) - \frac{1}{w(r(\theta_1), \theta_1)} \phi'(l(\theta_1)) \right) \end{aligned}$$

Note that the monopoly quality valuation correction term does not enter the optimal tax and subsidy because the monopoly quantity distortion is taken care of by the price subsidy in this implementation. The profits that the firm maximizes are exactly equivalent to \tilde{Y}^* , the socially valued output net of production costs.

Implementation with a prize mechanism:

The government can also simply purchase the innovation directly from the firm in exchange for a prize $G(\lambda, r)$ that depends on the step size (or, interchangeably, on the realized quality q) and on R&D investment. If the prize function is differentiable in its two arguments, the formulas for the marginal change in prize with respect to the step size or R&D investments can immediately be obtained by substituting for the wedges in the planner's first-order conditions, using the link between the wedges and the marginal prize with respect to product quality and R&D expenses.

$$\begin{aligned} s(\theta_i) &= \frac{\partial G(\lambda(r(\theta_i), l(\theta_i), \theta_i), r(\theta_i))}{\partial r(\theta_i)} + \frac{\partial G(\lambda(r(\theta_i), l(\theta_i), \theta_i), r(\theta_i))}{\partial \lambda} \frac{\partial \lambda(r(\theta_i), l(\theta_i), \theta_i)}{\partial r(\theta_i)} \\ \tau(\theta_i) \frac{\partial \pi(q(\theta_i), \bar{q})}{\partial q(\theta_i)} \frac{\partial \lambda(r(\theta_i), l(\theta_i), \theta_i)}{\partial l(\theta_i)} &= \frac{\partial G(\lambda(r(\theta_i), l(\theta_i), \theta_i), r(\theta_i))}{\partial \lambda} \frac{\partial \lambda(r(\theta_i), l(\theta_i), \theta_i)}{\partial l(\theta_i)} \end{aligned}$$

A.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 $C(k, \bar{q}) = \frac{k}{\bar{q}^\zeta}$, 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:

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

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

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

Profits are then given by

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

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

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

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 2, but also assume the special case in which $\rho_{\theta r} = \rho_{rl} = 1$, 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 $-\left(1 - \frac{1}{B_e}\right) - \frac{1}{B_e} \left(1 - \frac{1}{B_m}\right) < 0$ 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.

A.2 Proofs of the Propositions in the Main Text

Proof of Proposition 1:

Taking the first-order conditions of the planner's problem in (10) with respect to $l(\theta_i)$ and $r(\theta_i)$ for each $i = 1, 2$ and using the definitions of the wedges yields the formulas.

Proof of Proposition 2:

Taking the FOC of program P in (21) with respect to $r_t(\theta^t)$ yields:

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

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 (25).

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

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

Transform the derivative of the envelope condition:

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

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

Proof of Implementation Result:

For every period, define the following objects:

$$D_s(\theta^{s-1}, \theta_s) = E \left(\sum_{t=s}^{\infty} I_{(s),t} \left(\frac{1}{R} \right)^{t-s} \frac{\partial v_t}{\partial \theta_t} \middle| \theta^s \right)$$

$$Q_s(\theta^{s-1}, \theta_s) = \int_{\underline{\theta}}^{\theta_s} D_s(\theta^{s-1}, q) dq$$

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, $\frac{\partial v_t}{\partial \theta_t} = \phi'_t(l_t(\theta^t)) \frac{\frac{\partial \lambda(\theta^t)}{\partial \theta_t}}{\frac{\partial \lambda(\theta^t)}{\partial l_t}}$, 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:

$$T_t(\theta^t) = Q_t(\theta^{t-1}, \theta_t) - \frac{1}{R} E_t(Q_{t+1}(\theta^t, \theta_{t+1})) + \phi(l_t(\theta^t)) \quad (\text{A1})$$

Given the time separable utility and the assumption on the impulse response functions, the transfer hence depends on $\lambda_t, r_{t-1}, \theta_t$, 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_1 , 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_1 , 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 problem. By incentive compatibility, the firm will choose the allocation meant for its true type realization.

A.3 Computational Appendix

A.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 non-linear 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 ($2^{15} = 32768$) 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 et al. (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 et al. (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 the baseline value. These differences can represent actual difference 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 error 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.

A.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 (A1)). 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.

ONLINE APPENDIX [NOT FOR PUBLICATION]

OA.1 Compustat Data Matched to Patent Data

In this section, we redo our analysis on the sample made of only publicly traded firms, based on COMPUSTAT matched to patent data. For this purpose, we select our sample so as to make it as close as possible to the one in Bloom, Schankerman, and Van Reenen (2013). The sample selection procedure that follows Bloom, Schankerman, and Van Reenen (2013) keeps all firms who patent at least once since 1963, so that they can at least at some point be matched to the patent data (this is natural also in light of our theory, which focuses on innovating firms). The final unbalanced panel contains 736 firms that are observed at least four times in the period 1980 to 2001 and is essentially identical to the sample in Bloom, Schankerman, and Van Reenen (2013).⁴⁰ Table OA1 provides some summary statistics from the data.

TABLE OA1: SUMMARY STATISTICS IN THE COMPUSTAT AND PATENT DATA

Variable	Mean	Median
Sales (in mil. USD)	3133	494
Citations per patent	7.7	6
Patents per year	18.5	1
R&D spending / sales	0.043	0.014
Number of employees (000's)	18.4	3.8
Number of firms	736	

Note: The sample is selected to match as closely as possible the one in Bloom, Schankerman, and Van Reenen (2013), who keep firms that patent at least once since 1963 and which are observed for at least four years between 1980 and 2001.

⁴⁰The results are robust to this sample selection. We repeated the analysis on a much broader sample of 6,400 firms over the period 1976 to 2006 that could be matched to the patent data for any year (without restricting to firms that are observed for at least four years). The results on this alternative sample are similar and are available upon demand.

OA.2 Policies with a Finite Firm Life Cycle

OA.3 Robustness Checks on Parameters and Moments

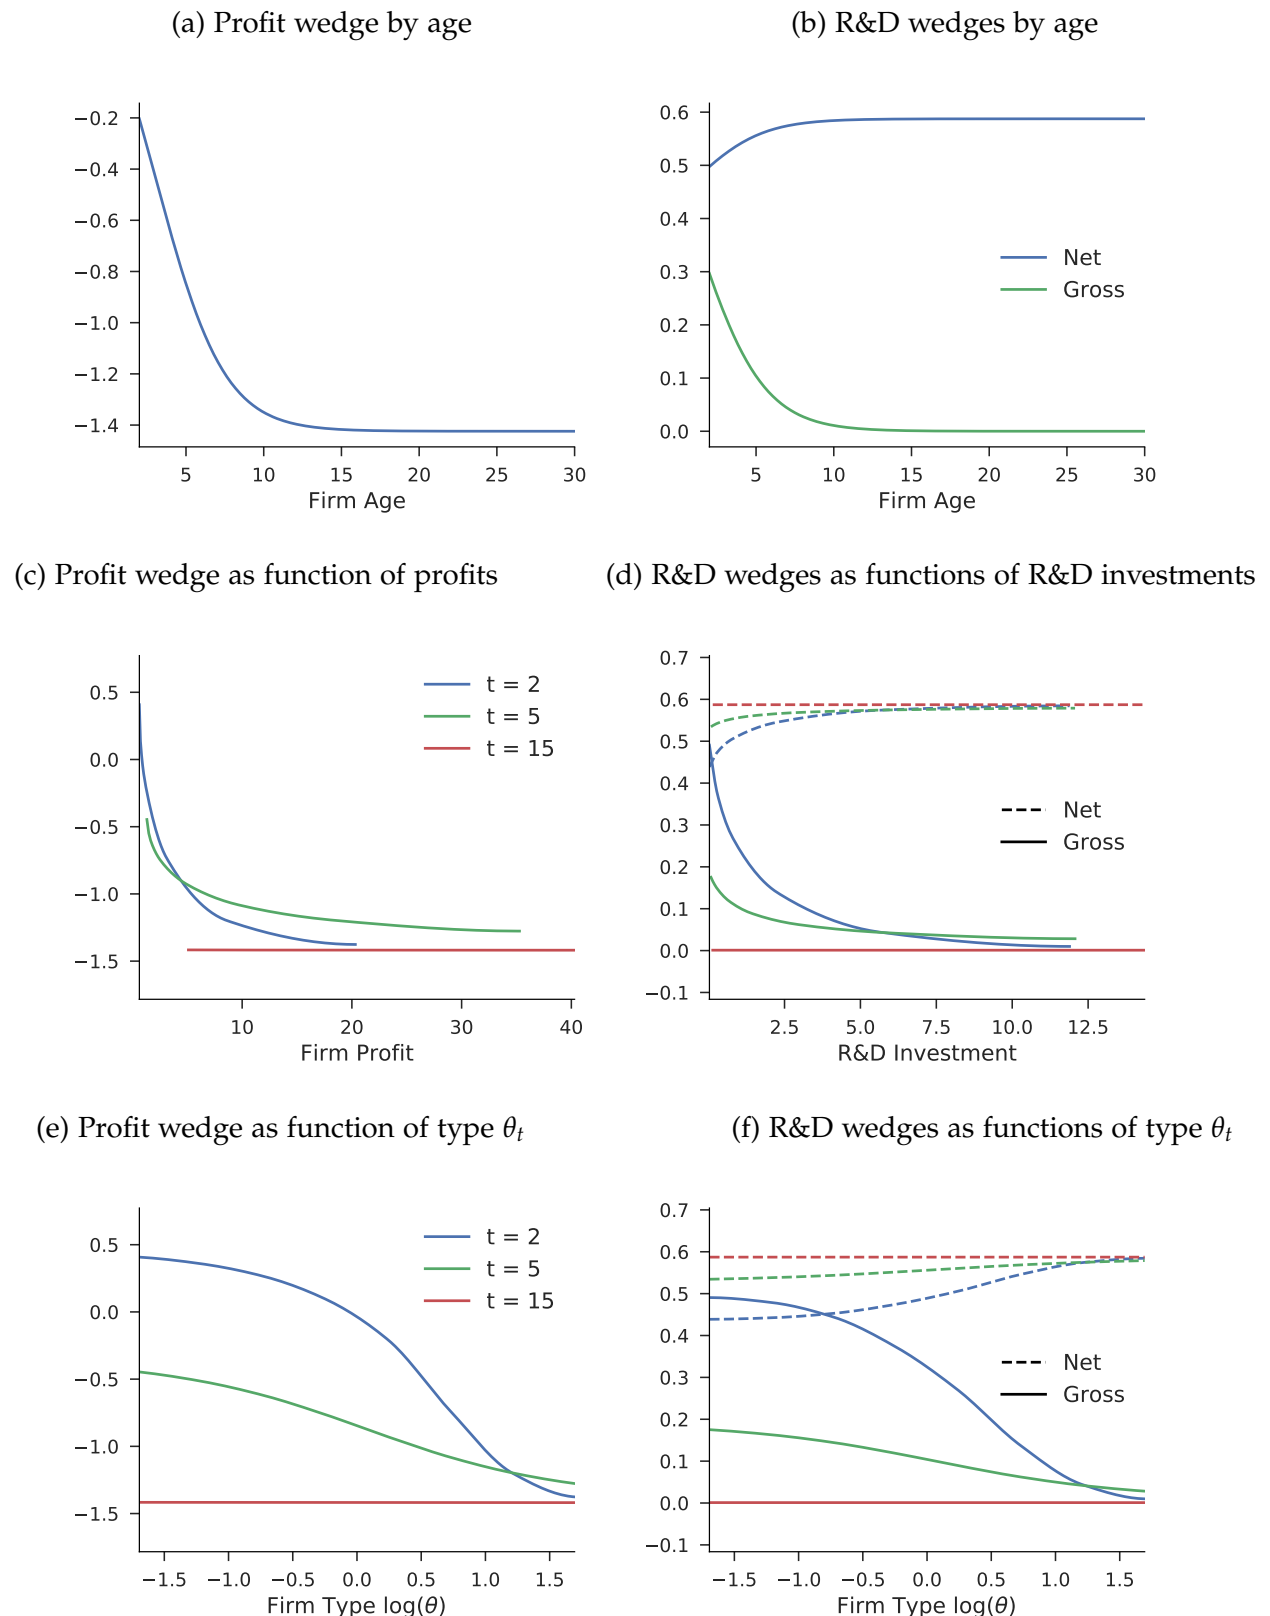
TABLE OA2: PARAMETER VALUES USING TWO-STEP GMM

Parameter	Symbol	Value
<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	0.48
R&D-type substitution	$\rho_{\theta r}$	1.84
Type variance	σ_ε	0.342
Type persistence	$\tilde{\rho}$	0.69
Scale of disutility	κ_l	0.72
Scale of R&D cost	κ_r	0.061
Effort cost elasticity	γ	0.94
Support width for θ_1	Θ^1	1.75
Production externality	ζ	0.018

TABLE OA3: MOMENTS USING TWO-STEP GMM

Moment	Target	Simulation	Standard Error
M1. Patent quality-R&D elasticity	0.88	0.96	(0.0009)
M2. R&D/Sales mean	0.041	0.034	(0.0025)
M3. Sales growth (DHS) mean	0.06	0.07	(0.005)
M4. Within-firm patent quality coeff of var	0.63	0.79	(0.0017)
Across-firm patent quality coeff of var:			
M5. Young firms	1.06	1.04	(0.0012)
M6. Older firms	0.99	0.89	(0.0016)
M7. Patent quality young/old	1.04	1.03	(0.0048)
M8. Spillover coefficient	0.191	0.190	(0.046)
M9. Elasticity of R&D investment to cost	- 0.35	-0.34	(0.101)

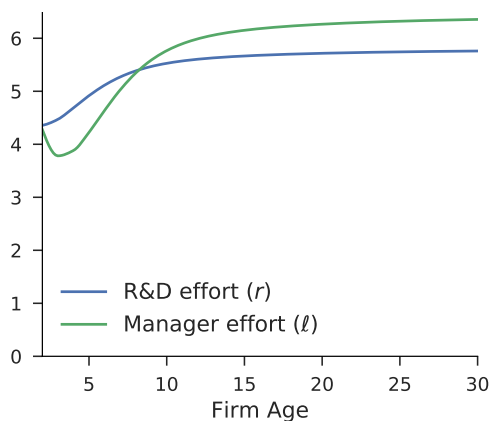
FIGURE OA7: OPTIMAL PROFIT AND R&D WEDGES FOR COMPUSTAT PUBLICLY TRADED FIRMS



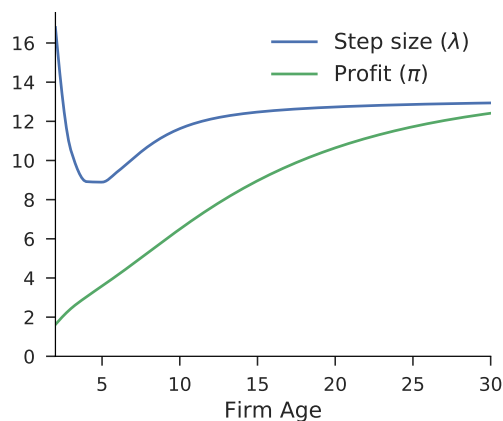
Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA8: OPTIMAL ALLOCATIONS FOR COMPUSTAT PUBLICLY TRADED FIRMS


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type

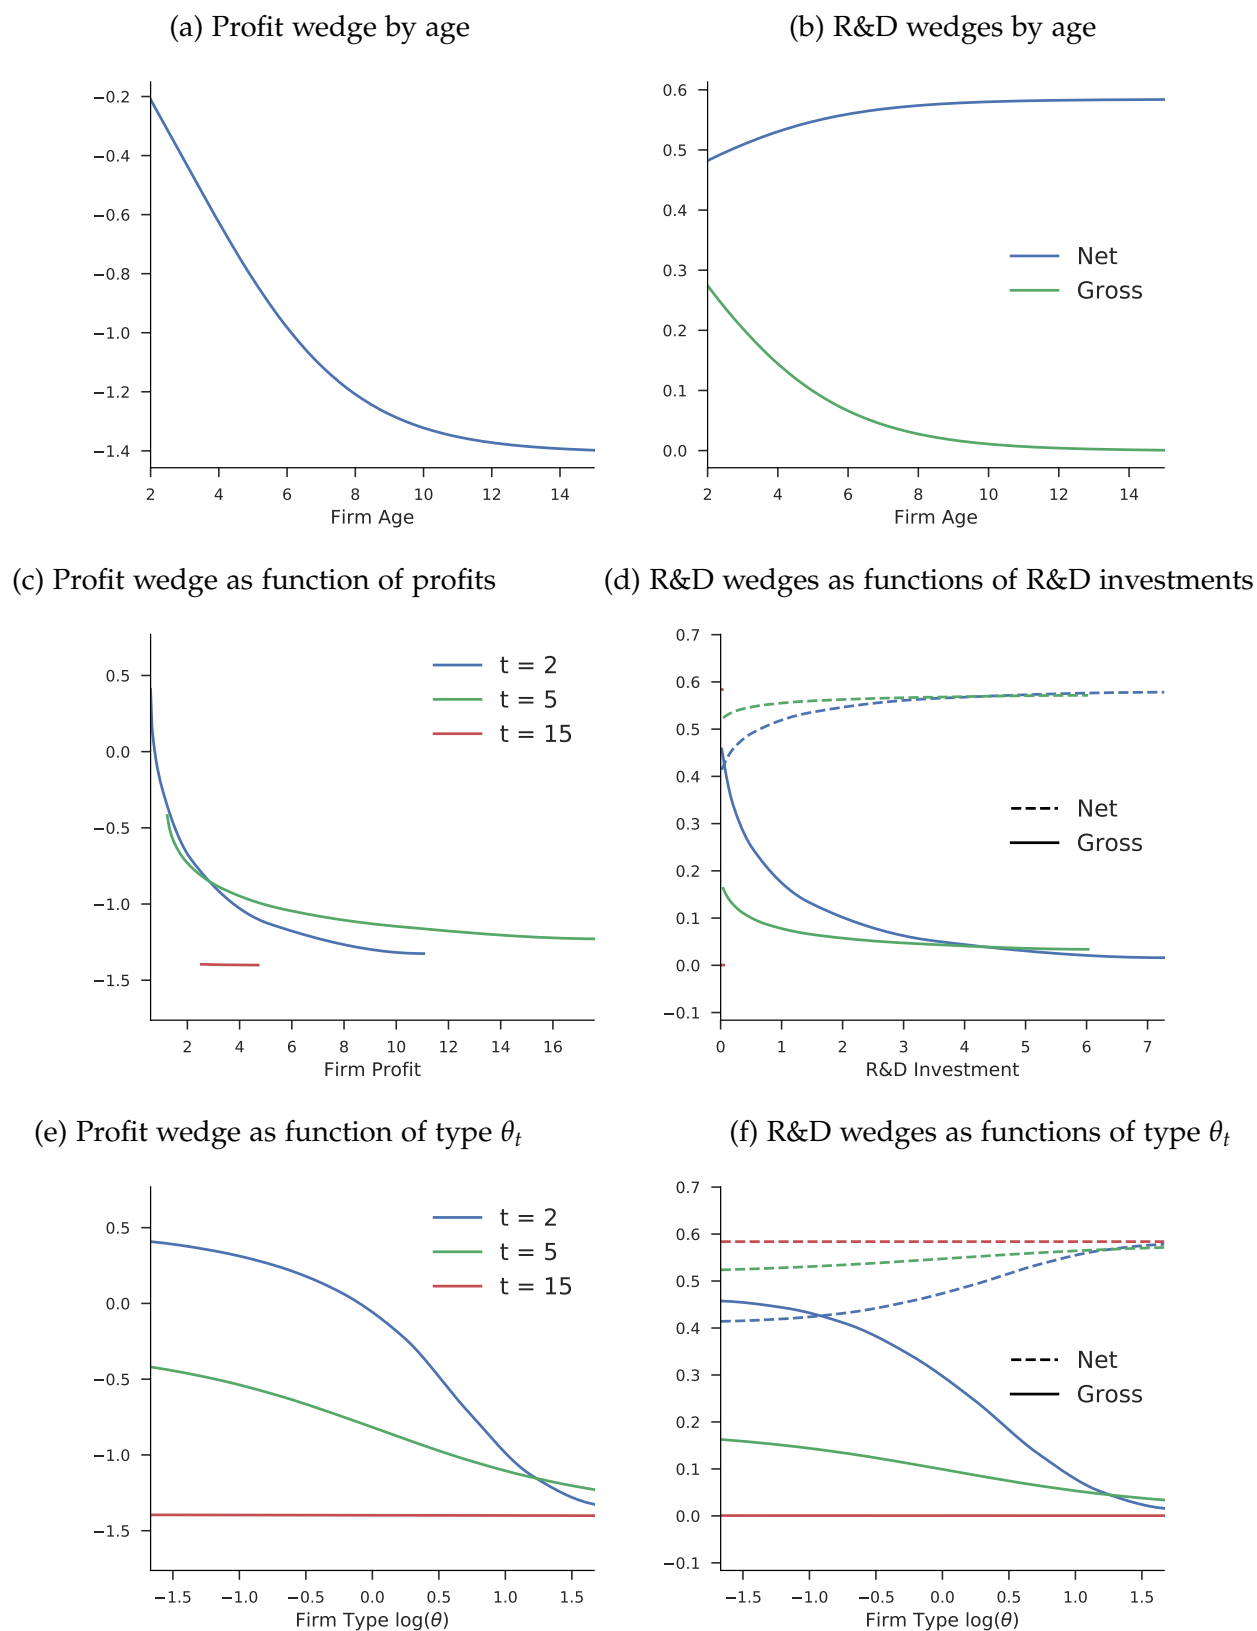


(d) R&D investments by Type



Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.


FIGURE OA9: OPTIMAL PROFIT AND R&D WEDGES FOR FINITE FIRM LIFE CYCLE $T = 15$




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA10: OPTIMAL ALLOCATIONS FOR FINITE FIRM LIFE CYCLE $T = 15$


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type




Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA11: OPTIMAL PROFIT AND R&D WEDGES USING TWO-STEP GMM


(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




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA12: OPTIMAL ALLOCATIONS USING TWO-STEP GMM


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type




Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA13: OPTIMAL PROFIT AND R&D WEDGES WITH AN AUTOREGRESSIVE PROCESS


(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




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA14: OPTIMAL ALLOCATIONS WITH AN AUTOREGRESSIVE PROCESS


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type



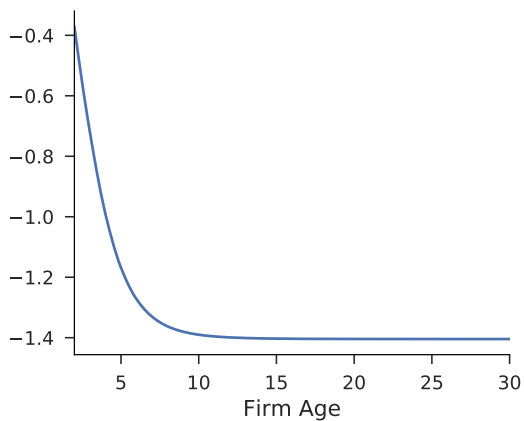
(d) R&D investments by Type




Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA15: OPTIMAL PROFIT AND R&D WEDGES WITH INCREASING PERSISTENCE p


(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




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA16: OPTIMAL ALLOCATIONS WITH INCREASING PERSISTENCE p


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type




Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA17: OPTIMAL PROFIT AND R&D WEDGES WITH PERSISTENCE $p = 0.5$


(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



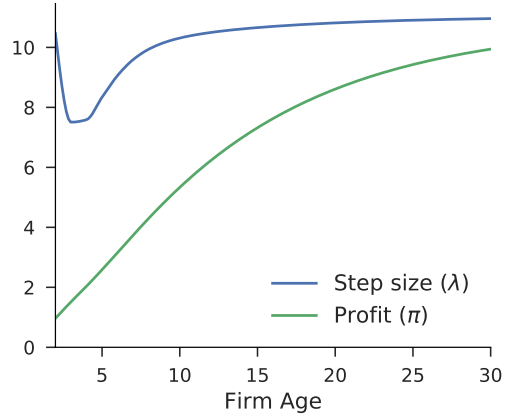
Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA18: OPTIMAL ALLOCATIONS WITH PERSISTENCE $p = 0.5$


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type



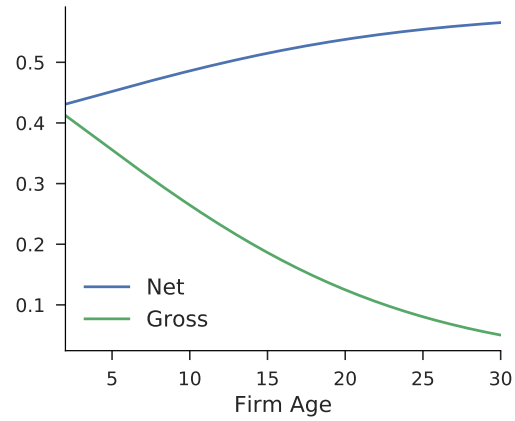
Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA19: OPTIMAL PROFIT AND R&D WEDGES WITH PERSISTENCE $p = 0.9$


(a) Profit wedge by age




(b) R&D wedges by age




(c) Profit wedge as function of profits



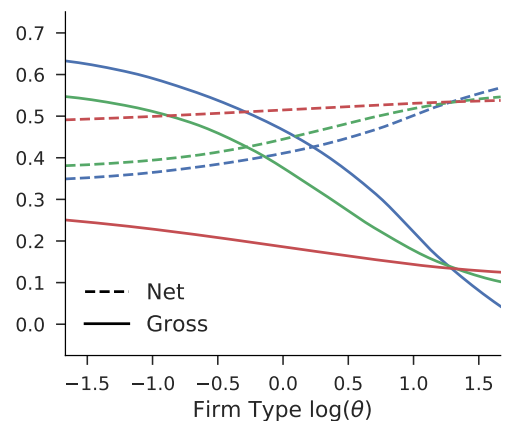
(d) R&D wedges as functions of R&D investments



(e) Profit wedge as function of type θ_t



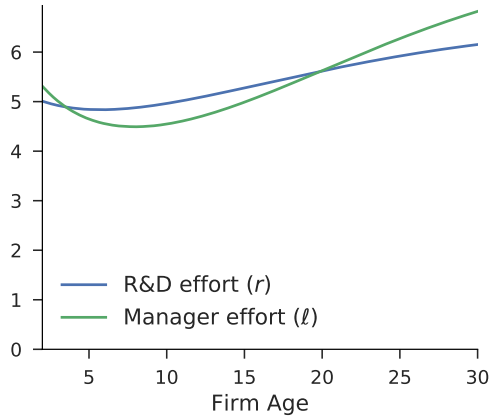
(f) R&D wedges as functions of type θ_t




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA20: OPTIMAL ALLOCATIONS WITH PERSISTENCE $p = 0.9$


(a) Investments and Effort by Age



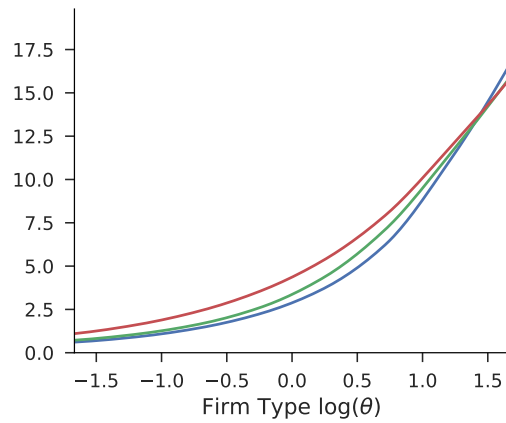
(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type




Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA21: OPTIMAL PROFIT AND R&D WEDGES FOR $\beta = 0.10$


(a) Profit wedge by age




(b) R&D wedges by age




(c) Profit wedge as function of profits



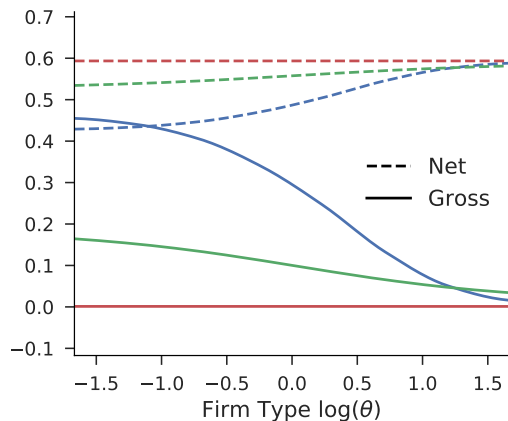
(d) R&D wedges as functions of R&D investments



(e) Profit wedge as function of type θ_t




(f) R&D wedges as functions of type θ_t




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA22: OPTIMAL ALLOCATIONS FOR $\beta = 0.10$


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type




Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA23: OPTIMAL PROFIT AND R&D WEDGES FOR $\beta = 0.25$


(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




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA24: OPTIMAL ALLOCATIONS FOR $\beta = 0.25$


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type




Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA25: OPTIMAL PROFIT AND R&D WEDGES FOR $\delta = 0.15$


(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




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA26: OPTIMAL ALLOCATIONS FOR $\delta = 0.15$


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type



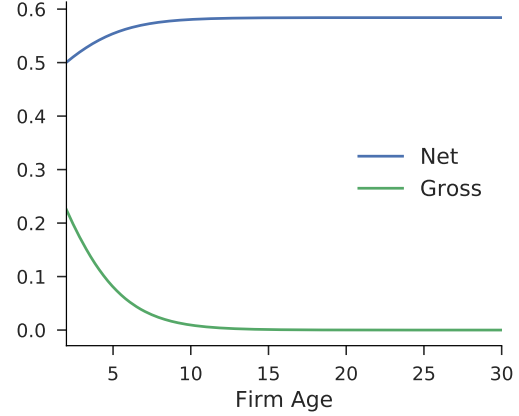
Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA27: OPTIMAL PROFIT AND R&D WEDGES FOR $\delta = 0.3$


(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




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA28: OPTIMAL ALLOCATIONS FOR $\delta = 0.3$


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type




(d) R&D investments by Type




Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.

FIGURE OA29: OPTIMAL PROFIT AND R&D WEDGES, OVERWEIGHTING MOMENT 1


(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




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA30: OPTIMAL ALLOCATIONS, OVERWEIGHTING MOMENT 1


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type

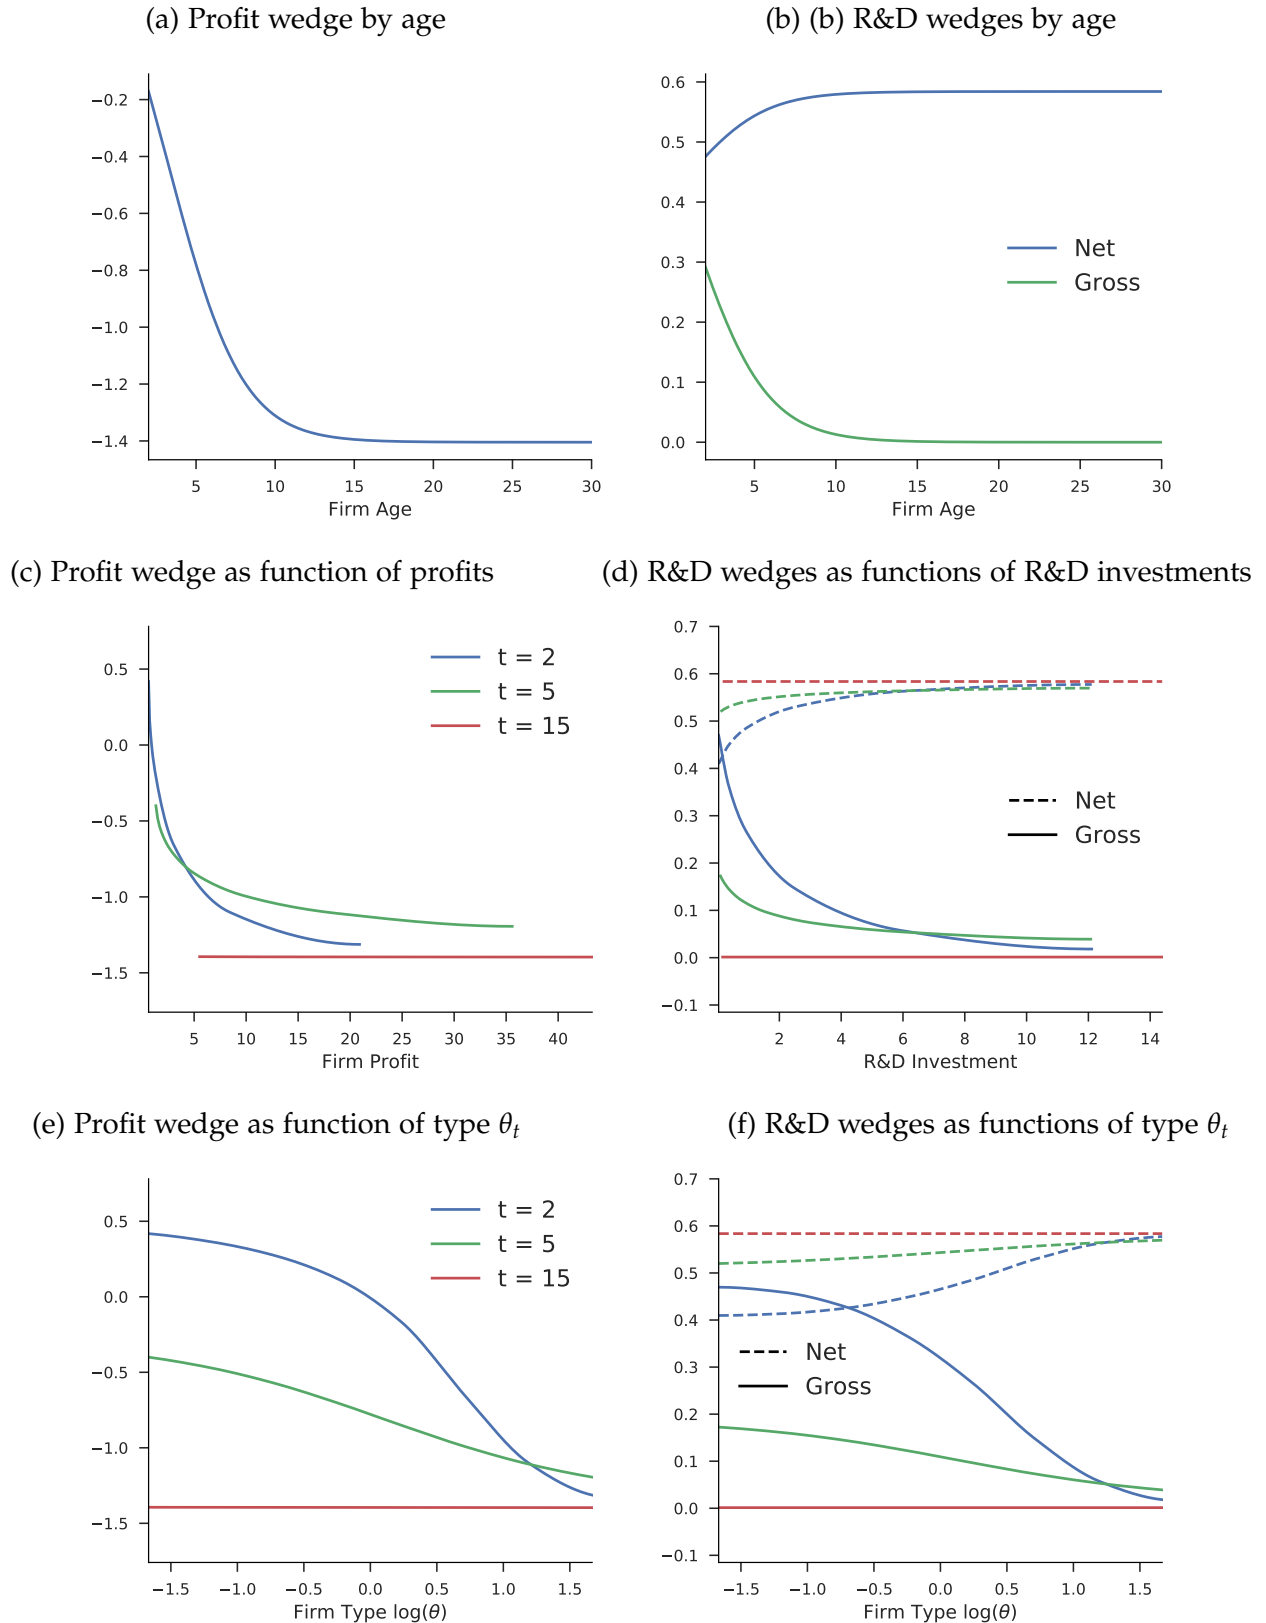


(d) R&D investments by Type



Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.


FIGURE OA31: OPTIMAL PROFIT AND R&D WEDGES, WITH $\eta = 1$




Notes: Panel (a) plot the average optimal profit wedge at different ages; Panel (b) plots the average optimal gross and net R&D wedges. Panels (c) and (d) plot, respectively, the optimal profit and R&D wedges for $t = 2, 5, 15$ for different levels of profits and R&D investments. Panels (e) and (f) plot the same wedges, but against firm productivity type θ_t .

FIGURE OA32: OPTIMAL ALLOCATIONS FOR $\eta = 1$


(a) Investments and Effort by Age




(b) Step size and Profits by Age



(c) Effort by Type



(d) R&D investments by Type



Notes: The figure depicts the optimal allocations for different ages and types of firms. Panel (a) shows optimal investments in R&D and effort for different ages; panel (b) shows the resulting step size and profits by age. Panels (c) and (d) depict, respectively, the optimal R&D effort and R&D investments for firms of different types for ages 2, 5, and 15.