



HCEO WORKING PAPER SERIES

Working Paper



HUMAN CAPITAL AND
ECONOMIC OPPORTUNITY
GLOBAL WORKING GROUP

The University of Chicago
1126 E. 59th Street Box 107
Chicago IL 60637

www.hceconomics.org

Optimal Taxation and R&D Policies*

Ufuk Akcigit Douglas Hanley Stefanie Stantcheva

November 28, 2016

Abstract

We study the optimal design of R&D policies and corporate taxation when the outputs of innovation are not appropriable in the absence of intellectual property rights policies and there are non-internalized technology spillovers across firms. 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. There is asymmetric information about firm productivity and about its stochastic evolution over time that prevents the first best solution to the technology spillover. The problem is thus posed as one of dynamic mechanism design with externalities. We characterize the optimal constrained efficient allocations over firms' life cycles and for firms of different productivities. We show that the constrained efficient allocations can be implemented either by a patent system plus a price subsidy for the monopolists' products, together with a parsimonious R&D subsidy function or, equivalently, by a prize mechanism. We estimate our model using firm-level data matched to patent data and quantify the optimal policies. Simpler innovation policies, such as linear R&D subsidies and linear profit taxes, lead to large revenue losses relative to the optimal mechanism.

Keywords: Corporate taxation, Tax credits, Subsidies, Innovation, Patents, R&D.

JEL Codes: H21, H23, H25, H32, O31, O32, O33, O38

*Akcigit: University of Chicago, NBER, and CEPR (e-mail: uakcigit@uchicago.edu); Hanley: University of Pittsburgh (e-mail: dohan@as.pitt.edu); Stantcheva: Harvard and NBER (e-mail: sstantcheva@fas.harvard.edu). We thank participants at the Society of Economic Dynamics Conference 2015, National Tax Association Meetings 2015, the American Economic Association meetings 2016 and numerous seminar participants at Chicago, Harvard, and Northwestern for feedback and comments. We thank Jessica Liu and Sanjay P. Misra 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

In a recent statement, Jason Furman, chairman of the Council of Economic Advisers (CEA) claimed that “the need to foster greater innovation and productivity growth is one of the most important economic challenges we face, and tax policy is one of several important levers that policymakers can use.”¹

Encouraging innovation is a key goal in many countries. Potential policies the government can implement to reach this goal are to improve competition, to regulate the intellectual property rights regime, to directly fund and perform R&D in public institutions, and to use tax policies that set the right incentives for firms to internalize some of the positive spillovers that their innovations have on society and on other firms. In this paper we focus on the latter and consider the optimal design of taxation and R&D policies. We combine elements from the macro innovation literature, theoretical tools of mechanism design, and firm-level data matched to patent data to discipline and quantify our analysis.

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, 50.56 billion USD on contracting with non Federally Funded Research and Development Centers (FFRDCs), as well as 27.8 billion USD on college and universities funding (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 reaches close to 25%.²

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 deduction, direct grants for research, contracting with private firms, subsidies for R&D costs, or direct funding in FFRDCs.³ The configurations of these many policies also vary widely.⁴ 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. Can we study the best set of policy tools for innovation endogenously, without restricting them *a priori*? What key parameters

¹“Encouraging Innovation and the Role of Tax Policy,” Remarks by Jason Furman, Chairman, Council of Economic Advisers. *Joint International Tax Policy Forum & Georgetown University Law Center Conference*, March 11, 2016. As prepared for delivery.

²Source: 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.

³They also directly fund areas perceived as being of high national priority, such as defense or health.

⁴For 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.

do optimal policies depend on? If the fully optimal policies are complex, how close could simpler policies come? In this paper, we attempt to answer these questions theoretically and quantitatively.

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. This is the idea of “building on the shoulders of giants.”⁵ 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, create a distortion as well, 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 returns to R&D are uncertain.

In a world without private information, the government could perfectly correct for the technology externality through a Pigouvian subsidy and for the inappropriability 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).

Empirically, 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. (2013a)). It would naturally be very difficult for the government to observe factors such as management quality, and to directly condition public policies on those factors, which is exactly the asymmetric information problem that is studied in this paper.

We focus on the optimal provision of incentives for innovation through the design of R&D policies and corporate taxation. Intellectual property rights (IPR) policies are distinct from, but intimately intertwined, with the latter. Indeed, if there is no IPR, any innovation immediately becomes public knowledge, profits are zero, and, regardless of the strength of subsidies for research inputs, there will be no investment in innovation. On the other hand, if there is full patent protection, part of the benefit from the subsidy for inputs is dissipated as monopoly profits. While our main concern

⁵The saying is traditionally attributed to Isaac Newton, from a letter to his rival Robert Hooke dated February 5, 1676: “What Descartes did was a good step. You have added much several ways [...]. If I have seen further it is by standing on the shoulders of giants.” It has been used extensively in the endogenous growth literature to talk about the technology spillovers.

is not the design of IPR policy, we do take IPR into account. In the first case we consider, we allow the government to freely set the optimal IPR jointly with the optimal R&D policy. In settings in which innovation quality is observable, the optimal IPR policy is very simple: it takes the form either of a prize system, or, equivalently, of a patent system combined with a product price subsidy that ensures efficient quantity. In the second case, we constrain the government to take the IPR as given, which leads to a partial optimum, and shows how optimal R&D policies should be set in the presence of (irremovable) patent protection.⁶

The model is presented in Section 2. 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 and start by posing the problem as a 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, and extend them by offering a new approach to allow for spillovers between agents (firms) in the presence of asymmetric information.

Our first contribution, in Section 5, is to solve for and characterize these second-best constrained efficient allocations with asymmetric information and spillovers. The optimal incentives for R&D trade-off a Pigouvian correction for the technology spillover and a correction for the monopoly distortion against the need to screen good firms from bad ones. We highlight that a crucial statistic is 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 productivity: the more complementary R&D investment is to firm productivity, the more rents a firm can extract if R&D investment are subsidized. This puts a brake on how well the government can set the Pigouvian correction and correct for the monopoly distortion., i.e., screening considerations can dampen the 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.

Another main result of our paper is that two parsimonious implementations of the optimal mechanism exist. First, the constrained efficient allocations can be implemented with a price subsidy on the monopolists' products that, for any given product quality, aligns the quantity produced with the socially optimal one, plus a comprehensive R&D subsidy that depends on firm age, current, one-period lagged, and first period quality, and current and one-period lagged R&D investment. Equivalently, it can be implemented by a prize mechanism that also depends on firms'

⁶Even if the quality of innovation is observable in our benchmark setting, except in Section 7, the first-best outcomes cannot be achieved because we want to channel investments to the most productive firms, and productivity is private information. While we do want the most productive firms to produce their innovations, we do not want them to appropriate the full surplus from it. IPR, when it can be freely set here, will achieve the efficient quantity conditional on quality, but will not achieve the efficient quality unless firm-productivity specific taxes are available (which is impossible with asymmetric information on firm type). Put differently, the problem solved in the mechanism design part of this paper (Sections 3 - 5) is about how to transfer surplus from firms to consumers without discouraging innovation.

innovation inputs (R&D investment) and the change in the product quality. We also discuss the age and cross-sectional patterns of the optimal policies, depending on the values of the parameters. If the fundamental technological factors do not systematically vary with age or with size, then the allocations of younger or lower productivity firms are more strongly distorted, i.e., they face higher marginal taxes on profits and higher R&D subsidies. This minimizes the informational rents extracted by high productivity firms.⁷

We take the model to the data in Section 6. We use COMPUSTAT firm 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, and growth rates and R&D intensities, among others. We use the technology spillover estimates in Bloom, Schankerman, and Van Reenen (2013) to discipline the magnitude of spillovers in our model. We then numerically simulate and quantify the optimal policies that we previously derived analytically. Younger firms on average face higher marginal taxes on profits and higher marginal R&D subsidies. The same also goes for lower productivity firms, conditional on age.

Our final contribution is to answer the key question of how close simpler innovation policies can come to approximating the full unrestricted mechanism. In Section 7, we compare the revenue raised from the full optimum to the one raised by a linear R&D subsidy combined with a linear profit tax, by linear age-dependent R&D subsidies and profit taxes, by a price subsidy on the product, or by a size-dependent R&D subsidy that varies with the level of R&D investments. In each case, the revenue losses are quite large. The nonlinearity of the optimal policies highlighted both analytically and numerically seems crucial in order to reap all the revenue gains.

Related Literature.

There is a long-standing 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 contributes to the new dynamic public finance literature that focuses on the dynamic income taxation of agents under idiosyncratic risk, including, among others, the papers by Farhi and Werning (2013), Golosov, Tsyvinski, and Werning (2006), Golosov, Tsyvinski, and Werquin (2014), and Werquin (2016). Most closely related in this literature are the papers by Stantcheva (2014) and Stantcheva (2016), which incorporate endogenous investments in human capital. Our paper considers the taxation of firms – rather than of individuals – when firms endogenously improve their productivity through R&D investments and there are spillovers across firms. Also related is the

⁷Higher subsidies do not imply that small or young firms invest more in R&D – in fact, they invest less, but require more incentives to do so.

paper by [Ales and Sleet \(2016\)](#), who study income taxation when firm productivity is endogenous because of manager talent (but not through the innovation channel as in our paper).

We also contribute to the theoretical and structural literatures on innovation, productivity, and growth. [Leahy and Neary \(1997\)](#) consider the effects of R&D subsidies in a model of oligopoly between firms. [Akcigit and Kerr \(2010\)](#) model heterogeneous innovations with spillovers like those in our model. [Acemoglu et al. \(2013\)](#) consider optimal industrial policies for the reallocation of resources in which firms are of heterogeneous qualities. [Atkeson and Burstein \(2015\)](#) study analytically and quantitatively the impact of taxes and subsidies on firms' investments in innovation. The key new elements in our paper relative to this literature are the asymmetric information in a dynamic setting with spillovers and the study of optimal R&D policies through a fully general direct revelation mechanism.

We build on the empirical literature on R&D and productivity to discipline our model and our empirical estimation.⁸ First, the empirical evidence on the importance of management practices ([Bloom and Van Reenen \(2007\)](#), [Bloom, Sadun, and Van Reenen \(2012\)](#), [Bloom et al. \(2013a\)](#)) lends support to the idea that firms are heterogeneous in their productivity 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 [Hall and Van Reenen \(2000\)](#), [Bloom, Griffith, and Van Reenen \(2002\)](#), [Bloom and Griffith \(2001\)](#), [Bloom, Chennells, Griffith, and Van Reenen \(2002\)](#), and [Griffith, Redding and Van Reenen \(2001\)](#). In the spirit of the present paper, [Serrano-Velarde \(2009\)](#) examines the heterogeneous impacts of R&D subsidies on firm investments. We directly make use of the technology spillover estimates from [Bloom, Schankerman, and Van Reenen \(2013\)](#).

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 on the value of the innovation, rather than on the firm's productivity to use research inputs. Worth mentioning, however, is a closely related and highly complementary paper by [Chari, Golosov, Tsyvinski \(2012\)](#) who focus on environments in which the value of the innovation (the quality in our model) is not known, but there is a market signal about it. In [Section 7](#), we also consider restricted policies for which the tax or subsidy rate do not condition on innovation quality, i.e., which do not assume that innovation quality is known to the planner.⁹

⁸Early papers that pioneered the use of patent data to study firms' innovation choices are [Pakes and Griliches \(1984\)](#), [Pakes \(1985\)](#), [Pakes \(1986\)](#), and [Pakes and Schankerman \(1986\)](#).

⁹Within the intellectual property policy protection literature, [Scotchmer \(1999\)](#) and [Cornelli and Schankerman \(1999\)](#) focus on the optimal patent renewal policy in a simple, stylized environment. Except in very special cases (see [Scotchmer \(1999\)](#)), the patent renewal policy is not the most general mechanism, but rather a restricted tool. [Kremer \(1998\)](#) studies patent buyouts as a way to foster innovation and counter the monopoly power that patents would provide. Information about the value of the patent is extracted through an auction which exploits the information of other firms in the same industry.

2 A Dynamic Innovation Model

Intermediate good producers' R&D decisions: The core of the model are the intermediate goods producers. Each intermediate good producer has product quality q_t at time t , that can be improved through R&D *investment* and R&D *effort*. Quality evolves according to:

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

where λ_t is the endogenous step size:

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

which depends on three components:

(i) *Observable R&D investments:* r_{t-1} denotes the resources that the firm invested in its 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 $M_t''(r_t) \geq 0$.¹⁰

(ii) *Unobservable R&D inputs or research effort:* Each firm also needs to provide some unobservable R&D inputs, which cannot be directly monitored by the government. The most natural interpretation is that unobservable investments relate to research effort, which is required in order to transform the material resources into an innovation output. The effort is denoted by l_t and requires a cost $\phi_t(l_t)$.

(iii) *Firm type:* Every firm has a type θ_t that determines the efficiency with which it produces quality (innovation) out of the given set of inputs r_{t-1} and l_t . The type is a potentially composite measure of the exogenous characteristics of a firm. For instance, θ may represent managerial productivity in innovation, an interpretation bolstered by recent literature on the importance and heterogeneity of management practices across firms (Bloom and Van Reenen (2007), Bloom, Sadun, and Van Reenen (2012), Bloom et al. (2013a)).

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:

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

¹⁰Taking a broad view of these material inputs is consistent with the fact that many types of material inputs and expenses are eligible for R&D tax credits or subsidies (Tyson and Linden, 2012).

so that higher realizations of productivity θ , higher R&D investments and 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).

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 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.

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 technology spillovers between firms. One firm's quality has a beneficial effect on the production costs of other firms. Such knowledge spillovers are one of the key elements in the endogenous growth literature, embodying the idea of "standing on the shoulders of giants" (see, among others, Klenow and Rodriguez-Clare (2005), Aghion and Howitt (1992)). They can reflect the direct use of better technologies and processes in production and learning from new technologies. The specific shape of the knowledge spillovers is as modeled by Acemoglu and Kerr (2010).¹¹

¹¹It would be easy, as an extension, to introduce additional knowledge spillovers in the step size production without changing the logic of the mechanism presented below. There would then be an additional Pigouvian correction term.

FIGURE 1: MODEL SUMMARY



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 produced competitively using the intermediate goods with the production technology:

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

where $Y(q_t(\theta^t), k_t(\theta^t))$ is the consumption of the good produced by producer θ^t as valued by consumers, $k_t(\theta^t)$ is the quantity produced, and $q_t(\theta^t)$ is the quality of the producer with history θ^t . We normalize the price of the final good to one.

Intellectual Property Rights (IPR) Regime and Market Specification: The IPR regime

here can be modeled through the structure of the market between the final and intermediate good producers. One way of capturing different IPR regimes is through different demand functions $p(q_t(\theta^t), k_t(\theta^t))$. We provide two examples here for illustration.

For instance, with full patent protection, the intermediate good producer has monopoly power in his intermediate good and faces a downward sloping demand curve derived from the final good's (competitive) market as a function of its quality and quantity, $p(q_t(\theta^t), k_t(\theta^t)) = \frac{\partial Y(q,k)}{\partial k}$.

Another example is a prize system, in which the government buys the innovation in exchange for a prize, and directly takes over the production. When the quality of the innovation is observable, the optimal prize system will lead to production at the socially efficient level, conditional on quality. There is a demand function that can mimic this prize mechanism: the efficient quantity (conditional on quality) can also be obtained through a price subsidy that inflates the intermediate good producer's valuation to equal the social one. The prize system's outcome can hence be mimicked by a demand schedule for the intermediate good given by $p(q, k) = \frac{Y(q,k)}{k}$.

Firm Life Cycle: Firms live for T periods (possibly, $T = \infty$) and can borrow at a gross 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 presented here (potentially with an IPR regime that generates a demand function), 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 :

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

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 $p(q, k)$ (which is specified according to the IPR regime in place as just explained), profits are:

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

Figure 1 summarizes the model in schematic form.

2.1 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 directly mandate production and R&D inputs for each type to maximize net household consumption, $W^{\text{first-best}}$, equal to total expected discounted output net of

production costs and R&D investment costs:

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

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.

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 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 exposition, we simplify the accumulation equation of quality to be

$$q_t = (1 - \delta)q_{t-1} + \lambda_t \quad \text{with} \quad 0 < \delta < 1 \quad (6)$$

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:

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

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

where the expectation operator is over histories θ^t . Finally, the full surplus can be extracted lump-sum from each producer of type θ^t and entirely transferred to households.

2.2 Asymmetric Information and Government Policies

2.2.1 Asymmetric Information

The core asymmetry of information, which holds throughout this paper, is that the history of type realizations θ^t and the unobservable R&D effort l_t are private information of each firm. In addition, other actions may or may not be observable and we consider several cases.

In the benchmark case, the government observes the full histories of the step size λ_t , the realized quality q_t , the R&D investment r_t , and the production k_t .¹² This amounts to the government being

¹²In the first two cases considered, product quality q_t is assumed to be observable. This is reasonable because, first,

essentially free to set the IPR policy, which, at the optimum, and because innovation quality is directly observable, is either a prize system or a patent system with a price subsidy that aligns the monopolist's valuation with the social one.

We also consider the case in which quantity $k(\theta^t)$ is unobservable as well, or, equivalently, cannot be conditioned on by the government. This means that the patent system has to be taken as given and only the R&D subsidy and corporate tax can be set to stimulate investment in innovation.

It is worth briefly discussing why, despite the fact that quality q_t is observable, the first best cannot be implemented in these two cases. Normally, since firms are risk-neutral, the technology externality could be corrected by a prize mechanism or through a patent system where the monopoly distortion is resolved through a price subsidy that inflates the monopolist's perceived value from a marginal unit produced to be equal to the social value. Under both of these mechanisms, all the surplus from innovation goes to the firm or manager, and the household has zero consumption. To extract some of that surplus from the firm and transfer it to households in a non-distortionary way, we would need a tax that depends directly on firm type, which is not feasible with asymmetric information. Any tax that depends on endogenous observables, such as quality, is going to be distortionary and move the manager's effort and investment decisions away from the socially optimal choice, pushing us into the mechanism design realm considered in this paper.

Under two sets of circumstances, asymmetric information would not be a problem: First, if managers and households get the exact same weight in the social objective, or, equivalently, if there is no revenue requirement for the government. Second, if the government does care about households but is able to recoup all the social surplus through a lump-sum tax over all firms (that does not depend on quality). In this case, some firms will end up paying more (potentially, much more) than their total revenues. This situation is hence infeasible if there is limited liability. Thus, the constraints that make asymmetric information binding are the presence of non-equal weights on households and managers (or, equivalently, revenue concerns) and limited liability for firms.

Finally, in Section 7, we consider cases in which quality is not observable. In this case, policy instruments are more restricted and take the form of taxes and subsidies that are linear (i.e., tax rates and subsidies do not depend on quality) or nonlinear in the inputs only. These policies do not require the government to observe the quality of innovations.¹³

there is a tight link between profits and quality in the model. Saying that the policy maker can condition on quality is equivalent to saying that he can condition on profits (or, sales). Second, several papers deal with the question of observing product (or innovation) quality. For instance [Cremer and Scotchmer \(1996\)](#) show that with correlation in the firms' private information on the innovation value, a mechanism can elicit this information without costs. [Kremer \(1998\)](#) also focuses on a way to elicit patent quality in auctions. Third, there are measures of quality such as patent citations that could be used (as we do in Section 6) albeit with some delay in time, to measure the quality of the innovation: the reward could be delayed until the measure of quality (e.g., citations) becomes available to the policy maker. This would not alter the theoretical results' main forces. Finally, a useful extension would be to study the case in which quality is observed with some noise and the planner can only condition on that noisy signal.

¹³With linear instruments, all that is required is that the government can see some measure of aggregate quality or profits and keeps a constant share of it (equal to the linear tax rate). There is no need to observe the identity of individual firms and individual quantities.

2.2.2 Direct versus Indirect Screening:

Why should the government screen agents indirectly through a mechanism rather than try to directly extract information on their type? Direct research funding such as through the National Science Foundation or the National Institute of Health often takes the direct screening approach: grant money is handed out after an involved screening process that tries to get at the “type” θ^t of the researcher or the project directly. There are several drawbacks to such direct screening. First, unlike decentralized instruments such as taxes and subsidies, these mechanisms are not easily scaleable. In addition, even if after detailed scrutiny, the government is able to get some signal about the difference between, say, SpaceX and Microsoft, it is not clear that this can be explicitly conditioned on in a systematic way in policy. Posting a menu of contracts and letting firms self-select seems more feasible. Second, a lot of discretion will be required by isolated institutions, which poses issues of accountability and transparency. This discretion also means that there may be a large temptation to extract all the surplus from high value innovations *ex post*, which would pose issues of commitment, and also leaves room for lobbying by interest groups. Third, unless the productivity of the potential grantees can be perfectly screened, grants which are not conditional on performance (i.e., on produced quality) but only on past performance or inputs will likely lead to underinvestment in the unobservable research inputs.

2.2.3 Government Policies Considered

We consider several types of government policies. The restrictions on the government’s tools assumed in each case have a close relation to the various informational structures described above. 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 and R&D effort. We do not constrain policy tools *ex ante*, but rather find the policy tools (prizes, price subsidies, or taxes) that will be able to implement these allocations.

We then turn to studying the shape of and revenue losses from restricted, Ramsey-type instruments which impose lower informational requirements on the government. In none of these cases does the quality of the innovation have to be observed by the government. We consider the optimal linear (i.e., size- and age-independent) corporate tax and R&D subsidies, the optimal linear, age-dependent corporate tax and R&D subsidies, and nonlinear R&D subsidies.

2.3 Discussion of the Model

We briefly discuss several of our modeling choices and their possible generalization.

First, the horizon of the problem is finite, but it could be very large or infinite. A finite and shorter horizon underscores the “firm life cycle” aspect of the optimal policies. One interpretation for T is as the horizon of policy making. It is realistic to think of policies as being set for a limited

time. For the exposition, it is assumed that firms disappear after period T . Alternatively, one could specify a terminal continuation value for each type of firm at time T which would be endogenous to policy.

Second, firms may be heterogeneous along many dimensions, such as their sector or the type of product. If the government or regulator wants to fine tune the policy for firms according to some observable vector of characteristics X , then the mechanism needs to condition on X . Since X is observable, this does not require adding any incentive constraints and only increases the state space to be kept track of.

Third, in the model firms make intensive margin decisions about how much to produce. Exit and entry can, however, be captured to some extent. The corner solution of zero production could represent exit (for some periods or forever after). The assumption needed is that, even if the firm is not active, its latent type θ continues to evolve according to its stochastic process, and that the underlying quality of the firm q_t remains in place (and possibly gets augmented by the step size if the firm still engages in R&D). It is reasonable to assume that the firm's "knowledge" stock q_t still remains in place and contributes to aggregate quality \bar{q}_t since that knowledge still exists and is known publicly. 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.

Fourth, we take the competition structure in the market as exogenously given. In reality, it may be endogenous to fiscal and innovation policy, a very interesting avenue for future research.

Finally, we consider two cases related to the ability of the government to control production by private producers: one in which quantity produced can be controlled and one in which it cannot. This allows us to discuss the impact of the IPR policy on the optimal policies. The case in which quantity is directly controlled by the planner corresponds to a situation in which the planner can choose the optimal IPR regime. In our setting, because quality of innovation 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). Equivalently, it corresponds to a case in which, even if there is a patent system granting monopoly power to the innovator, the planner can pay a nonlinear price subsidy to the monopolist that would align the private valuation of quantity with its social valuation.

The case in which quantity cannot be directly controlled mimics a situation in which the IPR policy, in the form of a patent system, has to be taken as given. This amounts to solving for the optimal profit tax and R&D subsidy given a patent system. The taxes and subsidies will then serve to indirectly counter the monopoly distortion created by patents.

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 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 .

The market structure between the intermediate goods and the final goods producer generates a demand function $p(q, k)$ for the intermediate goods, the shape of which depends on the IPR, in the exact same way as described in Section 2. 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 $(r(\theta_i), l(\theta_i), k(\theta_i), T(\theta_i))$ for $i = 1, 2$ and lets firms self-select allocations from this menu. We consider two cases.

Planning problem when quantity can be controlled: First, suppose that the planner can directly control quantity. The planner will then produce the socially efficient quantity conditional on a given quality. This is equivalent here to effectively optimizing the IPR policy, which, since quality is observable, takes the simple form of either a patent system plus a nonlinear price subsidy that aligns private and social valuations and solves the underprovision problem of the monopolist, or a prize system.

For any quality, the government should choose the socially optimal quantity, leading to output net of production costs $\tilde{Y}^*(q(\theta_i), \bar{q})$ for type θ_i .¹⁴ 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 the consumption of the households equal to:

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

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 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 2.1, and surplus is extracted in a lump-sum fashion from the

¹⁴This is because the optimal quantity to be produced is only conditional on quality and there is no reason to distort it (although the quality decision itself may be distorted).

firms, i.e.,¹⁵

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

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 constraint of type θ_2 and the participation constraint of type θ_1 will be binding.¹⁶ Indeed, at the first-best allocations and transfer levels, high 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 productivity firms, they could achieve a positive surplus by selecting the low productivity firm's first-best allocation. To prevent this from happening, the allocation of the low 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).

$$W = f_1 \left(\tilde{Y}^*(q_1(\theta_1), \bar{q}_1) - M(r(\theta_1)) - \phi(l(\theta_1)) \right) + f_2 \left(\tilde{Y}^*(q(\theta_2), \bar{q}) - M(r(\theta_2)) - \phi(l(\theta_2)) \right) - f_2 \left(\phi(l(\theta_1)) - \phi\left(\frac{w(r(\theta_1), \theta_1)l(\theta_1)}{w(r(\theta_1), \theta_2)}\right) \right) \quad (7)$$

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

¹⁵More precisely,

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

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

¹⁶As is usual in these types of screening problems, the slackness of the low type's omitted incentive constraint can be checked ex post.

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 firm'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 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.

Note that, while we need to take a stand on what the IPR in the laissez-faire is in order to define the wedges (although when quantity can be controlled, the IPR will itself be a choice parameter for the planner), this is just a matter of definition: the optimal allocations are unique and do not depend at all on how the wedges are defined since the latter are only used to intuitively characterize the allocations – it is only the benchmark relative to which the wedges can be interpreted as implicit taxes and subsidies that changes. It seems natural to define the implicit taxes and subsidies relative to a case in which profits are not zero, such as with patent protection, although this is not necessary: they measure the gap between marginal benefit and marginal cost, and the marginal benefit net of wedge could be zero, as long as the wedge is defined in absolute terms and not as a fraction of the marginal benefit. We could alternatively also define the wedges relative to the laissez-faire with a prize system. We explain below that this merely causes one term in the formula to drop out.

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 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 Productivity Firms.

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

$$s(\theta_1) = \underbrace{\left(f_1 \frac{\partial \tilde{Y}^*(q(\theta_1), \bar{q})}{\partial \bar{q}} + f_2 \frac{\partial \tilde{Y}^*(q(\theta_2), \bar{q})}{\partial \bar{q}} \right) \frac{\partial w(r(\theta_1), \theta_1)}{\partial r} l(\theta_1)}_{\text{Pigouvian correction}} + \underbrace{\left(\frac{\partial \tilde{Y}^*(q(\theta_1), \bar{q})}{\partial q} - \frac{\partial \pi(q(\theta_1), \bar{q})}{\partial q_1} \right) \frac{\partial w(r(\theta_1), \theta_1)}{\partial r} l(\theta_1)}_{\text{Monopoly quality valuation correction}} \\ + \underbrace{\frac{f_2}{f_1} \left(1 - \frac{\frac{\partial \log(w(r(\theta_1), \theta_2))}{\partial \log(r)}}{\frac{\partial \log(w(r(\theta_1), \theta_1))}{\partial \log(r)}} \right) \frac{\partial w(r(\theta_1), \theta_1)}{\partial r} l(\theta_1)}_{\text{Complementarity}} \underbrace{\phi' \left(\frac{w(r(\theta_1), \theta_1) l(\theta_1)}{w(r(\theta_1), \theta_2)} \right)}_{\text{Screening term}} \quad (8)$$

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

$$\tau(\theta_1) \frac{\partial \pi(q(\theta_1), \bar{q})}{\partial q(\theta_1)} = - \left(\frac{\partial \tilde{Y}^*(q(\theta_1), \bar{q})}{\partial q(\theta_1)} - \frac{\partial \pi(q(\theta_1), \bar{q})}{\partial q(\theta_1)} \right) - \left(f_1 \frac{\partial \tilde{Y}^*(q(\theta_1), \bar{q})}{\partial \bar{q}} + f_2 \frac{\partial \tilde{Y}^*(q(\theta_2), \bar{q})}{\partial \bar{q}} \right) \\ + \underbrace{\frac{f_2}{f_1} \left(\frac{1}{w(r(\theta_1), \theta_1)} \phi'(l(\theta_1)) - \frac{1}{w(r(\theta_1), \theta_2)} \phi' \left(\frac{w(r(\theta_1), \theta_1) l(\theta_1)}{w(r(\theta_1), \theta_2)} \right) \right)}_{\text{Screening term: Cost differential between high and low productivity firms}} \quad (9)$$

Proof. See Appendix A.2. □

The optimal implicit subsidy on R&D investment in (8) and the R&D effort wedge in (9) 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:* Recall that the wedge is defined as the implicit subsidy that would make the monopolist choose the planner's optimal allocation (starting from the 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 (the second term in each of the wedge formulas). If we had defined the wedge relative to a laissez-faire with a prize system, this term would drop out, as the monopoly distortion would be corrected directly through the prize system (or the equivalent price subsidy). The optimal R&D policies hence depend on the IPR policies, and several combinations of R&D policies and IPR policies could achieve the same outcome.¹⁷

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

¹⁷Accordingly, as we will see in the implementation below, when this monopoly valuation distortion is fixed with a price subsidy in the market for intermediate goods, this term disappears from the formula of the optimal marginal subsidy on R&D.

in the fraction of high productivity firms over low 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 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 productivity: if the elasticity is increasing in type, then R&D investments benefit disproportionately high 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.

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 productivity, as it then tightens the incentive constraint of the high 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 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 productivity firm’s R&D efforts depends on the difference in the marginal cost of producing the step size assigned to the low productivity firm between the low and the high productivity firm. The smaller this term the more tempting it is for the high productivity firm to imitate the low productivity one – this increases the optimal distortion on the low productivity firm’s R&D effort.

On the other hand, the high productivity firms’ allocations are set to their first-best levels since the low type’s incentive constraint is not binding, as shown in the following proposition:

Proposition 2. *Optimal Allocations for High Productivity Firms.*

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

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

Implementation. The constrained efficient allocations arising at the optimum of this direct revelation mechanism can be implemented in two ways.

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:

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

Appendix A.1 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.

Second, the government can 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 innovation quality and R&D expenses (see Appendix A.1). Given the straightforward map between the marginal grant and the wedges, all effects explained previously, such as the role of the relative return elasticities between the high and the low productivity firms or the difference in marginal effort costs apply directly to the marginal grant.

Planning problem when quantity can not be controlled: Next, suppose that the planner can no longer directly control quantity. This means that there is a monopoly distortion induced by the patent protection (that, by assumption, cannot be resolved with a price subsidy since we constrain the planner to not intervene in the market for intermediate goods). 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 planning problem, and hence the optimal wedges, are the same, but with $\tilde{Y}(q(\theta_i), \bar{q})$ replacing $\tilde{Y}^*(q(\theta_i), \bar{q})$ in (8) and (9).

Proposition 3. Optimal Allocations when Quantity Cannot be Controlled.

When quantity cannot be controlled, the optimal wedges in Propositions 1 and 2 apply with $\tilde{Y}(q(\theta_i), \bar{q})$ replacing $\tilde{Y}^(q(\theta_i), \bar{q})$ for $i = 1, 2$.*

When quantity cannot be directly controlled, the distortions in the R&D investment and effort are modified so as to indirectly compensate for the underprovision 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 $\frac{\partial \bar{Y}(q(\theta_i), \bar{q})}{\partial q(\theta)}$, is first-order and is proportional to the monopoly distortion, i.e., the gap between price and marginal cost.¹⁸ Hence, both the direct impact of R&D effort and investment on output and the indirect impact through the technology spillover will be amplified, making the R&D effort wedge smaller and the R&D investment wedge larger.¹⁹ In the numerical analysis in Section 6, it will be clear that the planner is able to induce less investment when the quantity is not controlled because this additional constraint makes incentive provision more costly.

4 A Direct Revelation Mechanism with Spillovers

We now return to the general 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 , the R&D investment r_t , and, depending on the case, the production k_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 $\sigma = \{\theta'_t(\theta^t)\}_{t=1}^T$. A reporting strategy generates a history of reports $\theta^{t'}(\theta^t)$. The government then assigns allocations as a function of the history of reports: $\{\lambda(\theta^{t'}), r(\theta^{t'}), k(\theta^{t'})\}$ and provides a transfer $T_t(\theta^{t'})$. Without loss of generality, we normalize the starting R&D investment for all agents to be $r(\theta^0) = r_0$.²⁰ 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 $\theta^{t'}$ (and, hence, had to invest $r(\theta^{t-1}(\theta^{t-1}))$ in the previous period and has to produce a step size of $\lambda_t(\theta^{t'}(\theta^t))$).

Suppose that the vector of aggregate qualities $\{\bar{q}_t\}_{t=1}^T$ is given. The continuation value after history θ^t under reporting strategy σ , denoted by $V^\sigma(\theta^t)$, is:²¹

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

$V^\sigma(\theta^t)$ depends on the report-contingent allocations specified by the government, $\{\lambda(\theta^s), r(\theta^s), k(\theta^s), T_s(\theta^s)\}_{s=1}^T$, although this dependence is implicit to lighten the notation.

¹⁸Formally, $\frac{\partial \bar{Y}(q(\theta_i), \bar{q})}{\partial q(\theta_i)} = \frac{\partial Y(q(\theta_i), k(q(\theta_i), \bar{q}))}{\partial q(\theta_i)} + \left(p(q(\theta_i), k(q(\theta_i), \bar{q})) - \frac{\partial C(k(q(\theta_i), \bar{q}), \bar{q})}{\partial k} \right) \frac{\partial k(q(\theta_i), \bar{q})}{\partial q(\theta_i)}$ where $k(q(\theta_i), \bar{q})$ is the quantity chosen to maximize profits by a monopolist with quality $q(\theta_i)$.

¹⁹Naturally, 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 when quantity cannot be controlled.

²⁰Since R&D investment r_0 is observable, we could always condition on it in the mechanism if it were heterogenous across firms.

²¹In sequential form the continuation utility as of the first period is:

$$V_1(\{\lambda(\theta^s), r(\theta^s), k(\theta^s), T_s(\theta^s)\}_{s=1}^T, \theta_1) = \sum_{t=1}^T \left(\frac{1}{R}\right)^{t-1} \cdot \left\{ \int_{\Theta^t} \{T_t(\theta^t) - \phi_t(l_t(\theta^t))\} P(\theta^t|\theta_1) d\theta_t \right\}$$

Let the continuation value under truthful reporting be $V(\theta^t)$. Incentive compatibility requires that, after every history, and for all reporting strategies σ :

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

In addition, limited liability or participation constraints require that, for all firms, $V(\theta^t) \geq 0$. Denote by \tilde{U} the realized lifetime utility given a sequence of realizations of the types, θ^T :²²

$$\tilde{U}(\theta^T) = \sum_{t=1}^T \left(\frac{1}{R}\right)^{t-1} \{T_t(\theta^t) - \phi_t(l_t(\theta^t))\} \quad (10)$$

4.1 A first-order approach

We use a first-order approach adapted from Pavan et al. (2014), which replaces all the incentive constraints of agents by their envelope conditions. 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). The integral form of this envelope condition at history θ^t is:

$$V(\theta^{t-1}, \theta_t) = \int_{\underline{\theta}}^{\theta_t} \frac{\partial V(\theta^{t-1}, m)}{\partial m} dm + V_t(\theta^{t-1}, \underline{\theta}) \quad (11)$$

This gives an expression for the informational rent forfeited to agent θ^t in period t to entice him to report his true type. 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 a Markov process):

$$I_{1,t}(\theta^t) = \prod_{s=1}^t \left(-\frac{\frac{\partial F^s(\theta_s|\theta_{s-1})}{\partial \theta_{s-1}}}{f^s(\theta_s|\theta_{s-1})} \right) \quad (12)$$

The impulse response function captures the persistence of the stochastic type process. For an AR(1) process such as $\theta_t = p\theta_{t-1} + \varepsilon_t$, the impulse response is simply:

$$I_{1,t}(\theta^t) = p^{t-1} \quad (13)$$

The envelope condition in its derivative form is given by:

$$\frac{\partial V(\theta^{t-1}, \theta_t)}{\partial \theta_t} = \mathbb{E} \left(\sum_{s=t}^T I_{t,s}(\theta^s) \frac{\partial \tilde{U}(\theta^T)}{\partial \theta_s} \right) \quad (14)$$

²²Implicitly, again this depends on the government's chosen allocation mechanism, i.e., $\tilde{U}(\theta^T) = \tilde{U}(\{\lambda(\theta^s), r(\theta^s), k(\theta^s), T_s(\theta^s)\}_{s=1}^T, \theta^T)$.

4.2 Planner's problem

The planner's objective is to maximize households' total consumption. Households own the firms and their consumption is equal to total output, minus production and R&D costs, and minus the transfers to the firms. We assume that the planner puts zero weight on the manager or the firm. This is easily relaxed: if the manager receives a weight of $(1 - k)$, all the wedge formulas below would be the same, but the screening term will be scaled by a factor $k < 1$. Intuitively, when there is a positive weight on the manager, the informational rent is not a pure loss from the planner's perspective, so it is optimal to not reduce it as much.²³ Suppose first that the sequence of aggregate qualities, $\bar{q} = \{\bar{q}_1, \dots, \bar{q}_T\}$ is fixed. We distinguish two cases.

Case 1: Quantity of production can be controlled. If the government can directly set the quantity to be produced, then for any $q_t(\theta^t)$ realized, the quantity of intermediate producer θ^t will then be chosen to maximize $Y(k_t(\theta^t), q_t(\theta^t)) - C(k_t(\theta^t), \bar{q}_t)$ which will yield $\tilde{Y}^*(q_t(\theta^t), \bar{q}_t)$, the socially optimal consumption minus production costs. We can hence directly substitute consumption net of production costs into the planner's objective and omit $k_t(\theta^t)$ as a control variable. The objective is:

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

Using the expression for \tilde{U} from (10), we can replace the sum of transfers $T_t(\theta^t)$ with \tilde{U} and the sequence of disutilities to obtain:

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

Using the expression for the informational rent that needs to be forfeited to each agent from (11), the expected discounted payoff to the planner is the "virtual surplus," i.e., the social surplus minus informational rents.

$$W(\bar{q}) = \mathbb{E} \left\{ \sum_{t=1}^T \left(\frac{1}{R} \right)^{t-1} \left\{ \tilde{Y}^*(q_t(\theta^t), \bar{q}_t) - M_t(r(\theta^t)) - \phi_t(l_t(\theta^t)) - V_1(\theta_1) - \frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t} \frac{\partial \tilde{U}_t}{\partial \theta_t} \right\} \right\}$$

Note that to satisfy the participation constraints, given the expression of the informational rents under the assumption that the first-order approach is valid, all that is needed is to set $V_1(\theta_1) = 0$. 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

²³As explained in Section 2, a knife-edge case is if the planner puts exactly the same weight on households and on the manager/firm, in which case $k = 0$ and there is no need for screening: the screening term drops out, the first-best corrective taxes and subsidies apply, all the surplus goes to managers/firms, and there is no attempt to transfer any of the surplus to the households. This can also be interpreted as a case in which the planner has no revenue requirement.

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 :

$$\int_{\Theta^t} q_t(\theta^t)P(\theta^t)d\theta^t = \bar{q}_t \quad (15)$$

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 $\{q_t(\theta^t), l_t(\theta^t), r_t(\theta^t)\}_{\Theta^t}$ so as to solve:

$$P(\bar{q}) = \max W(\bar{q}) \quad \text{s.t.:} \quad \int_{\Theta^t} q_t(\theta^t)P(\theta^t)d\theta^t = \bar{q}_t \quad \text{and} \quad q_t(\theta^t) = q_{t-1}(\theta^{t-1})(1-\delta) + \lambda(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t) \quad (16)$$

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

$$W(\bar{q}) = \sum_{t=1}^T \left(\frac{1}{R}\right)^{t-1} \left\{ \int_{\Theta^t} \{\tilde{Y}^*(q_t(\theta^t), \bar{q}_t) - M_t(r(\theta^t)) - \phi_t(l_t(\theta^t)) - \frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t} \left[\phi'(l_t(\theta^t)) \frac{\partial \lambda(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t) / \partial \theta_t}{\partial \lambda(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t) / \partial l_t} \right] - V_1(\underline{\theta}_1)\} P(\theta^t) d\theta^t \right\}$$

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.

$$P : \max_{\bar{q}} P(\bar{q}) \quad (17)$$

If quantity cannot be controlled: If the planner cannot directly choose the quantity, 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 planner's problem, denoted by $P^n(\bar{q})$, is the same as $P(\bar{q})$ above, replacing \tilde{Y}^* by \tilde{Y} , i.e., replacing $W(\bar{q})$ by:

$$W^n(\bar{q}) = \sum_{t=1}^T \left(\frac{1}{R}\right)^{t-1} \int_{\Theta^t} \{\tilde{Y}(q_t(\theta^t), \bar{q}_t) - M_t(r(\theta^t)) - \phi_t(l_t(\theta^t)) - \frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t} \left[\phi'(l_t(\theta^t)) \frac{\partial \lambda(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t) / \partial \theta_t}{\partial \lambda(l_t(\theta^t), r_{t-1}(\theta^{t-1}), \theta_t) / \partial l_t} \right] - V_1(\underline{\theta}_1)\} P(\theta^t) d\theta^t \quad (18)$$

Since the first-order approach is built on only necessary (but not necessary and sufficient) conditions, we perform a numerical *ex post* verification as in Farhi and Werning (2013) or Stantcheva (2016) to check that the candidate allocations are indeed incentive compatible.

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. We refer the reader to the detailed discussion in Section 3 about how the wedges are defined.²⁴

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

$$\tau(\theta^t) := \mathbb{E} \left(\sum_{s=t}^T \left(\frac{1-\delta}{R} \right)^{s-t} \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \frac{\partial \lambda_t(\theta^t)}{\partial l_t(\theta^t)} \right) - \phi'(l_t(\theta^t)) \right) \quad (19)$$

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

$$s(\theta^t) := M'_t(r_t(\theta^t)) - \frac{1}{R} \mathbb{E} \left(\sum_{s=t+1}^T \left(\frac{1-\delta}{R} \right)^{s-t-1} \frac{\partial \pi_s(q_s(\theta^s), \bar{q}_s)}{\partial q_s(\theta^s)} \frac{\partial \lambda_{t+1}(\theta^{t+1})}{\partial r_t(\theta^t)} \right) \right) \quad (20)$$

To simplify the notation, we use the following definitions.

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

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

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

²⁴As explained in Section 3, it makes sense – but is not necessary – to define the wedges relative to the laissez-faire economy in which firms receive profits from their innovations (i.e., there is a patent system or a prize system). We define them relative to the laissez-faire economy with a patent system, but could equally well define them relative to a prize system. The only difference, as in Section 3, is that the monopoly quality valuation correction term will then drop out of the optimal wedge formulas.

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.²⁵

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. The simple model in Section 3 already highlighted all the main intuitions which are here generalized.

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.,:

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

For instance, $\varepsilon_{l(1-\tau),t}$ is the elasticity of R&D effort to the net-of-tax rate $1 - \tau$.²⁶ 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 4. *Optimal corporate wedge and R&D wedge when quantity is controlled.*

(i) *The optimal profit wedge satisfies:*

$$\tau(\theta^t) = \underbrace{-\mathbb{E} \left(\sum_{s=t}^T (1-\delta)^{s-t} \eta_s \right) \frac{\partial \lambda_t}{\partial l_t}}_{\text{Pigouvian correction}} - \underbrace{\mathbb{E}(Q_t^*(\theta^t) - \Pi_t(\theta^t)) \frac{\partial \lambda_t}{\partial l_t}}_{\text{Monopoly quality valuation correction}} + \underbrace{\frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t}(\theta^t)}_{\text{Type distribution and persistence}} \underbrace{\frac{\phi'_t \lambda_{\theta t}}{\lambda_t} \left[\frac{1}{\varepsilon_{l,1-\tau}} \frac{1}{\varepsilon_{\lambda l,t}} + \rho_{\theta l,t} \right]}_{\text{Screening and incentive term Elasticity}} \quad (21)$$

²⁵Note 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).

²⁶Since there are no income effects for firms, the compensated and uncompensated elasticities are the same and equal to $\varepsilon_{l(1-\tau),t} = \frac{dl_t(\theta^t)}{d(1-\tau(\theta^t))} \frac{(1-\tau(\theta^t))}{l_t(\theta^t)} = \frac{\phi_{lt}}{(\phi_{ll,t} - \frac{\lambda_{ll,t} \phi_{lt}}{\lambda_{lt}})l}$.

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

$$\begin{aligned}
s(\theta^t) = & \underbrace{\mathbb{E} \left(\sum_{s=t+1}^T (1-\delta)^{s-t-1} \eta_s \frac{\partial \lambda(\theta^{t+1})}{\partial r_t} \right)}_{\text{Pigouvian correction}} + \underbrace{\mathbb{E} \left((Q_{t+1}^*(\theta^{t+1}) - \Pi_{t+1}(\theta^{t+1})) \frac{\partial \lambda(\theta^{t+1})}{\partial r_t} \right)}_{\text{Monopoly quality valuation correction}} \\
& + \frac{1}{R} \mathbb{E} \left(\underbrace{\frac{1 - F^1(\theta_1)}{f^1(\theta_1)} I_{1,t+1}(\theta^{t+1}) \phi'_{t+1}(l(\theta^{t+1}))}_{\text{Type distribution and persistence}} \frac{\lambda_\theta \lambda_r}{\lambda \lambda_l} \underbrace{(\rho_{lr} - \rho_{\theta r})}_{\text{Relative complementarity}} \right)
\end{aligned} \tag{22}$$

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

$$\int_{\Theta^t} \frac{\partial \tilde{Y}^*(q_t(\theta^t), \bar{q}_t)}{\partial \bar{q}_t} P(\theta^t) d\theta^t = \eta_t \tag{23}$$

Proof. See Appendix A.2. □

The optimal wedges in (21) and (22) 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 he would if he were maximizing profits, this term decreases the profit wedge and increases the R&D wedge. If we had defined the wedge relative to the laissez-faire with a prize system (or, equivalently, with a patent system and a subsidy on intermediate goods prices), this term would drop out. When the government can freely optimize IPR policy, he can decide how much to do through the intellectual property design and how much to do through the taxes and subsidies.²⁷

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 (21) 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 (22) is increasing in the expected effect of R&D investments on the next period's step size $\frac{\partial \lambda_{t+1}}{\partial r_t}$.

Screening considerations may, however, go in the opposite direction.

²⁷As we will also see in the implementation, when the monopoly distortion can be resolved through a price subsidy that, at given aggregate quality, aligns the social and private valuation of quantity and makes the monopolist produce the socially optimal quantity, the marginal profit tax or R&D subsidy no longer depend on the monopoly distortion.

3) Screening and incentives: The screening term arises because of the asymmetric information. We discuss three effects at play here.

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 only 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 p , $I_{1,t} = p^{t-1}$ so that the impulse response is fully determined by the persistence parameter.

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 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.²⁸ 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 ($\varepsilon_{l,1-\tau,t} \varepsilon_{\lambda,t}$). The informational rent inefficiency increases in the complementarity of effort to firm productivity $\rho_{\theta l,t}$. Recall from the simple illustration in Section 3 that the effort wedge on the low productivity firm was higher when the high 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 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 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 on the step sizes of high productivity firms (at any given effort level), which makes it easier for them to mimic

²⁸This is naturally reminiscent of the inverse elasticity rule in Ramsey taxation.

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.

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.²⁹ 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.³⁰ 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 Some special cases

If there is no asymmetric information or if the planner weighs households with the exact same weight as the manager, the screening terms drop out. The optimal profit wedge and the optimal R&D subsidy are then equal to the Pigouvian and monopoly quality valuation corrections, as in Section 2.1. Externalities are corrected under full information (and tailored to each productivity history θ^t), and the surplus is extracted from firms in a lump-sum manner,

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. On the other hand, if types are *iid*, then the screening term is zero after the first period: since agents are risk-neutral and types exhibit no persistence, there is no gain in spreading the screening distortion over different periods. As a result, in all periods after period 1, the optimal R&D subsidy and profit wedge are equal to the Pigouvian and monopoly quality valuation corrections.

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). Hence, R&D investment incentives are only provided to counter the technological spillovers and the monopoly distortion.

²⁹This interpretation was not possible in the simple one-period model in Section 3, where there was no uncertainty.

³⁰Recall 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 .

5.1.2 Age profile of optimal policies

The optimal policies will generically change over time. Since we consider a cohort of firms, time patterns are equivalent to age patterns. Age patterns can come from three sources, which are conceptually very different.

The first reason for time dependent policies is the finite horizon, which 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 . This age-pattern can be dampened or fully eliminated by letting the horizon go to infinity or by stipulating a terminal value for each firm as a function of the period T investments. 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 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.

The second, perhaps more fundamental, reason is that the state-contingent policy is set at time 1 under full commitment from the planner. As a result, it is the distance to period 1 that induces age patterns. The optimal corporate wedge and R&D wedge decline with age, as long as the impulse response is below 1 (as is the case for instance with an AR(1) process with persistence parameter $p < 1$). 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. Hence, the smaller effective informational rents warrant less distortion in the allocations. The age patterns induced by the logic of the screening of firms will be less visible for a shorter horizon. As argued above, it may be reasonable to view T as the horizon of the planning problem, after which policies are no longer committed to.

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. The quantitative analysis in Section 6 will shed light on them.

5.1.3 Cross-sectional profile of optimal policies

When thinking of the cross-sectional patterns of the optimal wedges, it is important to bear in mind that a higher R&D wedge does not mean a higher investment in R&D; and, similarly, 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 productivity firms invest much less than high productivity firms and this pattern is not overturned despite the incentive provision.

From the formulas, it is clear that higher 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.

The comparative statics of the monopoly valuation correction term and the screening term with respect to firm productivity are ambiguous. Among others, they depend on the shape of the hazard rate, $\frac{1-F^1(\theta_1)}{f^1(\theta_1)}$ and the impulse response function. This ambiguity hence carries over to the optimal wedges.

Quantitatively, in Section 6 we will see that the wedges will be declining with firm type. This does fit with the screening logic of the simple two type illustration in Section 3 where it was the low productivity firm's allocation which was distorted in order to reduce the informational rents of the high productivity type (whose allocation was set at the efficient level).

5.2 Optimal allocations conditional on the patent system

When quantity cannot be directly controlled, it is as if the government can no longer optimize the IPR policy and has to take the patent system as given, without being able to directly fix the monopoly distortion through a price subsidy or a prize mechanism. Each firm directly chooses its quantity to maximize profits, and, at given quality and aggregate quality, the output net of production costs for the producer of history θ^t is $\tilde{Y}(q_t(\theta^t), \bar{q}_t)$. Let

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

be the marginal impact of quality on \tilde{Y} . Accordingly, the planner solves program $P^n(\bar{q})$ above, yielding the same optimal wedges as in Proposition 4, with Q_t^* replaced by Q_t and $\tilde{Y}^*(q_t(\theta^t), \bar{q}_t)$ replaced by $\tilde{Y}(q_t(\theta^t), \bar{q}_t)$.

There are two differences relative to the wedges in the case in which quantity can be controlled, (21) and (22), which are driven by this term substitution.

When quantity is chosen by the intermediate goods producer in the private market, it is set to maximize profits and not social surplus. In this case, 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.³¹ This monopoly quantity correction term is positive and always makes the profit wedge smaller and R&D subsidy larger. This is intuitive: the higher the monopoly

³¹Formally, $\frac{\partial \tilde{Y}(q_t(\theta^t), \bar{q}_t)}{\partial q} = \frac{\partial Y(q_t(\theta^t), k_t(q_t(\theta^t), \bar{q}_t))}{\partial q_t(\theta^t)} + \left(p(q_t(\theta^t), k_t(q_t(\theta^t), \bar{q}_t)) - \frac{\partial C(k_t(q_t(\theta^t), \bar{q}_t), \bar{q}_t)}{\partial k} \right) \frac{\partial k_t(q_t(\theta^t), \bar{q}_t)}{\partial q_t(\theta^t)}$ where $k_t(q_t(\theta^t), \bar{q}_t)$ is the quantity chosen to maximize profits by a monopolist with quality $q_t(\theta^t)$.

distortion, and the less the monopolist internalizes the social value from an increase in quality. The same amplification will appear in the Pigouvian correction, as there is an additional effect on social welfare through the adjustment in quantity: when aggregate quality increases, quantity produced increases, which has a first-order positive effect on social welfare. The positive externality is in that sense also amplified.

Naturally, not being able to control quantity directly represents an additional constraint, and, hence, a cost on the planner. In Section 6 it will be clear that just being able to control quantity (e.g., through the price subsidy described above) will lead to large welfare gains. When quantity cannot be controlled the effort and R&D investments of the firms are much lower, despite higher R&D wedges and lower (even negative) effort wedges.

5.3 Extensions: Different types of research and qualities

Different types of observable R&D investments: Suppose that there are in fact 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 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 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 (22) 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 productivity, $\rho_{\theta l}^j - \rho_{\theta r}^j$.

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 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 of a firm, 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. \(2014\)](#), 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 (22) then tells us that R&D investments with the highest spillovers (highest $\eta_t^j = \int_{\Theta^t} \frac{\partial Y^*(q_t(\theta^t), \bar{q}_t^1, \dots, \bar{q}_t^J)}{\partial \bar{q}_t^j} P(\theta^t) d\theta^t$) are the ones that should be subsidized most (bearing in mind that their complementarities with effort and firm productivity may dampen the benefits from subsidizing them).

5.4 Implementation through Taxes, Subsidies, and Prizes

Decentralizing the constrained efficient allocation. Until 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 $T(q^T, k^T, r^T)$ that depends on the full sequence of all observables, i.e., on the history of quality q^T (or, interchangeably, step size λ^T), quantity k^T (in the case when quantity is assumed observable), and R&D investment r^T , 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. The immediate next question is what the most parsimonious such implementing tax function is.

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.³² However, the shape and level of the tax function that implements the constrained

³²For 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 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.

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 in fact be much more 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). This generalizes the implementation explained in detail in the simplified model of Section 3.

The constrained efficient allocation from program $P(\bar{q})$ (when quantity is observable) can then be implemented in two ways, which from a theoretical point of view are equivalent. The first implementation features a price subsidy $s_p(k, q)$ such that the post-subsidy price perceived by the intermediate good producer is $p(k, q)(1 + s_p(k, q)) = \frac{Y(k, q)}{k}$. In this case, the private producer will maximize profits equal to $Y(k, q) - C(k, \bar{q})$ conditional on q , which is exactly the social surplus from production k . This price subsidy should be combined 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 .

Second, the government could set up a prize mechanism, through which he purchases the new innovation flow (i.e., the step size) λ_t from the firm in each period, and produces the socially optimal quantity of the good of quality $q_t = (1 - \delta)q_{t-1} + \lambda_t$. Here, the government becomes the central owner of the intellectual property and keeps adding to his stock every period, in exchange for a prize. The prize amount $G_t(\lambda_t, r_t, r_{t-1}, q_1)$ paid for an innovation λ_t depends on firm age, current and lagged R&D investments, and the initial quality q_1 .

The constrained efficient allocation from program $P^n(\bar{q})$ (when quantity cannot be controlled) is implemented with a similar comprehensive tax function, but without the price subsidy, as there is no intervention at all in the private market between final goods and intermediate goods producers.

Proposition 5. Implementation Results.

Assume that the impulse response functions $I_{1,t}(\theta^t)$ only depend on θ_1 and θ_t for all t .

(i) The constrained efficient allocation from program $P(\bar{q})$ (when quantity can be controlled) can be implemented with a price subsidy on the intermediate good, such that the post-subsidy price perceived by the intermediate good producer is $p(k, q)(1 + s_p(k, q)) = \frac{Y(k, q)}{k}$, and a comprehensive age-dependent tax function $T_t(q_t, r_t, q_{t-1}, r_{t-1}, q_1)$ that depends on the first period quality q_1 , current and lagged quality q_t and q_{t-1} , and current and lagged R&D, r_t and r_{t-1} .

efficient allocation would be different.

(ii) It can also be implemented by a prize mechanism, in which the government purchases the new innovation from firms in exchange of a prize $G_t(\lambda_t, r_t, r_{t-1}, q_1)$, and produces the socially optimal quantity.

(ii) The constrained efficient allocation from program $P^n(\bar{q})$ (when quantity can not be controlled) can be implemented with a comprehensive age-dependent tax function $T_t^n(q_t, r_t, q_{t-1}, r_{t-1}, q_1)$.

Proof. See Appendix. □

6 Quantitative Investigation

The theory presented in Sections 2 to 5 is not dependent on functional form assumptions and we were able to describe the comparative statics on the main primitives. In this section, we provide empirical content to the theoretical model, by estimating it and numerically illustrating the optimal policies.

6.1 Data and Summary Statistics

We use firm-level accounting data from COMPUSTAT data matched to the U.S. Patent and Trademark Office (USPTO) patent data from the NBER data website (as described in detail in [Hall et al. \(2001\)](#)), containing over three million patents with their forward citations (see also [Jaffe and Trajtenberg \(2002\)](#)).

Among all firms that can be matched between the patent and COMPUSTAT datasets, we select our sample so as to make it as close as possible to the one in [Bloom, Schankerman, and Van Reenen \(2013\)](#). These authors do a very careful job in estimating precise technology spillovers, which we will use for our estimation. In addition, our theory highlights some new important moments from the data, related to firm productivity, heterogeneity, and uncertainty, which have not yet been the focus of the earlier innovation literature and which we need to compute. For consistency, we compute these moments on the same sample as the one from which the technology spillover has been estimated 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 715 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 \(2013\)](#).

The 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.

Table 1 provides some summary statistics from the data. The large heterogeneity of firms in the data, as captured by the spread between mean and median is striking. The median firm has

sales of 494 million USD, while the mean firm has sales of 3,133 million USD. The ratio of R&D over sales is very skewed: while the median is at 1.4%, the mean is 4.3%. The same goes for the innovation process: the number of patents a firm receives per year has a median of 1 and a mean of 18.5.

TABLE 1: 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.

Map between the model and the data: The great advantage of the patent data matched to COMPUSTAT 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 depreciated stock of citations per patent to date $q_t = (1 - \delta)q_{t-1} + \lambda_t$. Profits and sales can be directly measured.

6.2 Estimation

To estimate the model presented in Section 2, we start by parameterizing it. The functional forms used are described in Table 2. 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 3. 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 4. 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: We parameterize our model as summarized in Table 2.

TABLE 2: 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	$\phi_t(l_t)$	$\kappa l_t^{\frac{1+\gamma}{1+\eta}}$
Cost of R&D	$M_t(r_t)$	$\kappa r_t^{\frac{1+\gamma}{1+\eta}}$
Stochastic type process	$f^t(\theta_t \theta_{t-1})$	$\log \theta_t = p \log \theta_{t-1} + (1 - p)\mu_\theta + \epsilon_t$
Distribution of heterogeneity θ_1	$f^1(\theta_1)$	$f^1(\theta_1) = \frac{I_{\Theta^1}(\theta_1)}{\theta_1[\underline{\theta}_1 - \bar{\theta}_1]}$
Initial quality level	q_0	0

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

The cost function decreases in aggregate quality \bar{q}_t , and the strength of the externality is measured by ζ . The step size is multiplicatively separable in labor 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, $\rho_{\theta l} = \rho_{\theta r} = 1$. Given that the sign of $\rho_{l r} - \rho_{\theta r}$ determines the sign of the screening term in the optimal R&D subsidy (as shown in Proposition 4), 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, with elasticities of, respectively, $\frac{1}{\gamma}$ and $\frac{1}{\eta}$. Finally, the stochastic process for firm productivity type is a geometric random walk, with persistence p . The shock ϵ_t follows a normal distribution with mean zero and variance σ_ϵ . The parameters are either calibrated externally or estimated internally to match the key features of the data. This is summarized in Table 3. Appendix A.2 solves for the optimal quantities, prices, and profits of the model with this parameterization.

Externally Calibrated Parameters: The effort elasticity is set to match a Frisch elasticity of 0.5 as in Chetty (2012), which implies $\gamma = 1$. The profit parameter β is set to 0.15, a typical value as discussed in Guner et al. (2008). The exponent on the R&D cost function, η , is set as in Akcigit

TABLE 3: PARAMETER VALUES

Parameter	Symbol	Value
<i>External Calibration</i>		
Effort cost elasticity	γ	1
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
<i>Internal Calibration</i>		
R&D share	α	0.390
R&D-type substitution	$\rho_{\theta r}$	0.861
Type variance	σ_ϵ	0.253
Type persistence	p	0.71
Scale of disutility	κ_l	0.88
Scale of R&D cost	κ_r	0.048
Support width for θ_1	Θ^1	1.98
Production externality	ζ	0.022

TABLE 4: MOMENTS

Moment	Simulation	Target
M1. Patent quality-R&D elasticity	0.50	0.57
M2. R&D/Sales median	0.014	0.013
M3. Sales growth (DHS) mean	0.08	0.074
M4. Within-firm patent quality coeff of var	0.67	0.77
Across-firm patent quality coeff of var:		
M5. Young firms	1.17	1.10
M6. Older firms	0.71	0.63
M7. Patent quality young/old	2.00	1.88
M8. Spillover coefficient	0.191	0.188

and Kerr (2010). The depreciation δ and the (long-run) interest rate R are standard. The average level of productivity is normalized to $\mu_\theta = 0$, while the initial R&D stock is normalized to $r_0 = 1$.

Moments and Identification: Table 4 lists the data moments that we match. The second column provides the value of the moment in the simulations, and the third column gives the target value of each moment in the data. In this section, we discuss the identification of the parameters of the model.

Let the vector of the eight endogenously estimated parameters be denoted by $\chi = (\alpha, \rho_{\theta r}, \sigma_\epsilon, p, \kappa_l, \kappa_r, \zeta, \Theta^1)$.

We chose the parameters so as to minimize the loss function:

$$L(\chi) = \sum_{k=1}^8 \left(\frac{\text{moment}_k^{\text{model}}(\chi) - \text{moment}_k^{\text{data}}}{\text{moment}_k^{\text{data}}} \right)^2$$

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.

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. In Appendix A.3 we provide a Jacobian matrix that reports the sensitivity of each moment to each parameter. This way, we verify, at least locally, that the moments that we use to identify certain parameters are indeed informative.

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). We use the value provided by Bloom, Schankerman, and Van Reenen (2013) as a target. Not surprisingly, this moment informs the complementarity (or elasticity of substitution) parameter $\rho_{\theta r}$ in the innovation production function (see Table 2).

R&D Intensity, M2: The second moment is the median ratio of R&D spending to firm sales, which is a measure of the R&D intensity of a firm. It is computed directly in our sample and matches the values in earlier papers based on U.S. firm data very closely. The level of 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. This intuition is verified in the Jacobian matrix in Appendix A.3.

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. 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 p in our model.

Across-firm Patent Quality Variation by Age, M5-M6: The fifth and sixth moments capture the variation in quality across firms. This cross-sectional variability measure gauges the degree of heterogeneity across firms and is computed separately for young firms and older firms. Young firms are defined as those who just enter the COMPUSTAT sample (i.e., just become publicly traded).

They do typically appear before that for some time in the patent data, where we can track them.³³ As is intuitive, these moments are mainly determined by the dispersion of the type distribution 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. Again, young firms are taken at their entry point in the COMPUSTAT sample and the stock of their patents and citations at that point is taken from their years to date in the patent data. As the horizon gets shorter and the benefits from R&D effort are reduced, firms attempt to save on the cost of R&D effort. The extent to which they reduce their R&D effort depends on the marginal cost of effort. Therefore this particular moment mainly informs the estimation about the scale parameter of the cost of R&D effort κ_l .

Spillover Coefficient, M8: Finally, 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 patent quality of a firm on the depreciated R&D stock weighted by the extent of technological proximity of other firms in the economy. They instrument for this R&D stock 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 setting the net cost of R&D for each simulation and generating simulated economies. As described in the Computational Appendix, this is achieved by exogenously shocking the scale parameter κ_r . We then regress the patent quality in the model on the R&D stock 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 .

6.3 Results

Having estimated the parameters of the model, we can simulate the optimal allocations and wedges, presented in analytical form in Section 5.

To facilitate an interpretation of the wedges as tax and subsidy rates, we slightly redefine the profit and R&D wedges, respectively, as fractions of profits and R&D costs. The R&D subsidy rate $\tilde{s}(\theta^t)$ is now the fraction of the cost $M(r)$ that the firm does not have to pay, while the profit wedge $\tilde{\tau}(\theta^t)$ is the fraction of profits that the firm pays. In addition, to mimic more closely the linear taxes and subsidies explored in Section 7, we define the R&D wedge as the gap between marginal costs and marginal benefits of R&D, taking into account the R&D effort wedge, i.e., $\tilde{s}(\theta^t)$ measures the gap relative to a laissez-faire with patent protection, but taking into account that there is simultaneously a tax $\tilde{\tau}(\theta^t)$ on profits. More precisely, several graphs below depict in addition to

³³We tried a similar estimation using firms less than 3 years old and less than 5 years old in the COMPUSTAT sample, with extremely similar results.

$s(\theta^t)$.³⁴

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

For the case in which quantity can be controlled, we show the wedges that apply once the optimal price subsidy is in place: hence, the wedges measure the additional distortion that needs to be imposed at the optimum in addition to the price subsidy. Thus, the monopoly quality valuation correction term does not enter the wedge formula (recall the discussion in Section 5 about this term dropping out when the wedges were defined relative to a laissez-faire with a prize system, i.e., when quantity is chosen optimally conditional on quality).

Gross incentives and net incentives: A brief discussion of gross and net incentives for R&D is useful here. 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 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 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. The net incentive is driven by the difference between the gross linear subsidy \tilde{s} and the tax τ : $s = \tilde{s} - \tau$. 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 Age Patterns of the Optimal Allocations

As explained in Section 5, age patterns can arise for three reasons, the finite horizon, the logic of the screening problem in which policies are set at time 1 with full commitment, and the age-dependency

³⁴They are related to the wedges from Proposition 4 through

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

of the primitives of the model. Recall also that a higher profit wedge represents a larger implicit tax on firm profits, while a higher R&D wedge represents a larger implicit subsidy on R&D expenses.

Figure 2 plots 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 different ages, averaged over firm type at a given age.

Younger firms simultaneously have their profits taxed at a higher rate and their R&D investment expenses subsidized at a higher rate. The profit wedge falls from around 50% to -10%. The R&D gross wedge falls from around 50% to 0%, while the net R&D subsidy falls from 14% to 11.3%. Thus, the net incentive for R&D does not decline as dramatically.

The logic of the screening, explained in Section 5, is at play here. Types are less than fully persistent (the estimated persistence parameter is 0.71) so the screening terms in Proposition 4 are largest early in life when the firm has the most potential informational rent 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 term.³⁵

Figure 3 plots the optimal allocations as a function of firm age. The left panel depicts the optimal inputs, R&D expenses r , and R&D effort l , while the right panel depicts the step size λ and the profits π . The paths of 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, it is the screening considerations which dominate, while in the latter part, it is the 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. Focusing on R&D effort first, this is because the R&D effort wedge decays over time and hence R&D effort increases more. After the mid-point of the life cycle, it is the finite horizon considerations that dominate, making it less worthwhile to put in effort and investment given the shortening horizon left to reap the benefits. As explained in Section 5, this life cycle consideration could be eliminated by making the horizon much longer or infinite or specifying a non-zero terminal value, increasing in the terminal quality q_T .

The pattern of R&D expenses is driven by the path of R&D effort. Although the R&D subsidy declines over time, R&D investments are initially increasing because their return is increasing in the amount of effort provided. Given the hump-shaped path of inputs, profits and step size follow the same pattern.³⁶ Profits are increasing initially as the quality stock is built over time.

6.3.2 Cross-sectional Patterns of the Optimal Allocations

Figure 4 plots 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 productivities for ages $t = 2, 5, 10$, and 20. Lower

³⁵The Pigouvian correction term itself is also declining due to the life cycle consideration, but that effect would disappear with an infinite horizon, unlike the decay of the screening term, which would still occur.

³⁶The initial dip in the step size happens because we calibrate the period 1 productivity distribution to match the heterogeneity in the data, while in the other periods the type distribution is obtained from the persistence implied by the data.

productivity firms face higher profit wedges and higher R&D wedges, and the wedges decline monotonically in firm type. The trade-off between reducing informational rents and distorting allocations is again at play: a larger distortion in the allocation of a given type can reduce the informational rent given to other types. Higher wedges for lower productivity firms means that lower productivity firms' allocations are more distorted relative to those of high productivity firms, which helps reduce the informational rent for higher productivity firms.

Figure 5 shows the optimal inputs for firms of different qualities for these same ages. Higher productivity firms should optimally provide more effort and invest more in R&D. Given 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.³⁷

6.3.3 Optimal Allocations and Wedges When Quantity Cannot be Controlled

When quantity cannot be controlled, there is no price subsidy that can align the monopolist's quantity choice with the socially optimal one. As a result, the R&D effort and investment wedges pick up the monopoly quality valuation term correction and are, respectively, smaller and larger to provide more incentives for the firm to provide innovation inputs.

Figure 6 depicts the optimal profit and R&D wedges averaged across firms at each age. Figure 8 shows the wedges as a function of firm productivity at different ages. For both figures, the decreasing pattern is driven by the same screening considerations explained in Section 6.3.1 for the case in which quantity could be controlled by a price subsidy. The levels are however different: the profit or effort wedge is lower and consistently negative, providing a net production subsidy at the margin. The R&D wedge is consistently higher. Both are due to the need to stimulate firms' investment in quality to make up for their lower valuation of it because of the (here, irremovable) monopoly distortion.

Figure 7 depicts the average inputs and outputs at each age, while Figure 9 depicts the inputs and outputs as a function of firm type for different periods. It is clear that, despite the more generous incentives provided, the fact that quantity cannot be controlled represents an additional costly constraint on the planner. The inputs and outputs that can be requested from firms are much lower than in the case where quantity can be controlled. Since there is underinvestment to start with, this is not a desirable outcome.

The welfare gains from just offsetting the monopoly power through a price subsidy explored numerically in Section 7 are large, emphasizing that the monopoly distortion creates a big loss in efficiency. Indeed, the welfare loss from the unrestricted mechanism when quantity cannot be controlled relative to the one where quantity can be controlled is 73.6%. In this framework, leaving the patent system unaffected while subsidizing research inputs generates substantially lower

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 productivity firms already invest much less than high productivity firms and this pattern is not overturned despite the incentive provision.

benefits than when the patent system can be replaced by a prize system as described in one of the implementations in Section 5.4.

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

We now numerically illustrate the comparative statics that were described already for the theoretical formulas in Proposition 4.

The persistence of the firm’s productivity process affects the optimal policies very significantly. Figures 10 and 11 depict, respectively, the optimal wedges and the optimal allocations for a higher value of the persistence ($p = 0.9$) and a lower value of the persistence ($p = 0.5$). Recall that the persistence estimated in the data falls in between these two values, at 0.71. With a low persistence, both wedges are smaller after period $t \geq 2$ because the screening term is dampened, and quickly decays to zero (the wedges for age 10 and age 20 become identical as the screening term has completely faded and only the Pigouvian correction remains).

With a high persistence, the allocations are tilted so that in all periods, higher productivity firms provide more R&D effort and R&D investment, while lower productivity ones provide less. With lower persistence, wedges and allocations are much flatter, especially in later periods.

Figures 12 and 13 depict the optimal wedges and the allocations for the special case $\rho_{\theta r} = \rho_{lr}$ and the case in which the complementarity between R&D and type is higher ($\rho_{\theta r} = 1.2 > \rho_{lr}$). In the former case, recall that R&D investment does not affect the firms’ incentive constraints and, hence, the optimal wedge is only equal to the Pigouvian correction. When $\rho_{\theta r}$ is larger, R&D mostly benefits higher productivity firms and tightens their incentive constraints: the optimal R&D wedge is then smaller. This is especially true for earlier periods when the screening term – which is driven by the complementarity coefficient – has not yet decayed. As expected, when $\rho_{\theta r}$ is larger, it is costlier to provide incentives for R&D so the inputs of R&D effort and R&D investment, the step sizes, and profits are smaller, especially for high productivity firms (the informational rent gains from larger R&D investments of whom would be prohibitively costly).

Finally, Figures 14 and 15 depict the wedges and allocations for a weaker ($\zeta = 0.01$) and a stronger technology spillover ($\zeta = 0.03$). Unsurprisingly, a larger spillover leads to lower R&D effort wedges, higher R&D investment wedges, much higher inputs in R&D effort and R&D investment, and, accordingly, higher profits and step sizes.

7 Simpler Innovation Policies

Until now we have considered a fully unrestricted mechanism that does not place constraints on the policy tools available to the government, except to subject allocations to incentive compatibility due to the asymmetric information.

In this section, we consider restricted, simpler policies, and ask two questions: First, what is

the optimal simpler policy in each restricted class of policies considered? Second, how large are the revenue losses from these restricted simpler policies relative to the full mechanism? The policies considered in this section do not require the government to be able to observe innovation quality.³⁸

7.1 Optimal Simpler Policies and Revenue Losses

We solve for the optimal policies within each restricted class numerically, using the same estimated parameters as for the unrestricted mechanism from Section 5. We then compute the revenue shortfall relative to the revenue obtained in the unrestricted mechanism.

In line with the discussion in Section 6, there are two ways to specify the R&D subsidy rate, depending on whether or not R&D expenses are deductible from corporate income subject to the profit tax. In this section, we report both the gross subsidy rate \tilde{s} (the subsidy rate if R&D expenses are not tax deductible) and the net subsidy rate s .

Linear age-independent policies. We start with the simplest possible specification, which is a linear R&D subsidy s and a linear tax on profits, τ . Finding the optimal linear policies in this restricted class yields a profit tax of 31.4% and a R&D subsidy of 40.5%. The effective subsidy on R&D, as explained right above, is much smaller at 9.1% = 40% – 31%. The revenue loss relative to the full mechanism is very large, namely 65%. These policies are depicted in Panel A of Table 5. Both the profit tax and the (gross) R&D subsidy are larger than their current effective tax rates and subsidies in the U.S..

Linear age-dependent policies. We then consider age-dependent, but still linear, policies, for which firms pay different linear profit taxes and receive different subsidies based on their age group. The first policy, in panel B, allows for two age brackets: 1-15 years old firms and 16-30 years old firms. The younger age group pays higher profit taxes and receives a larger R&D subsidy, in line with the optimal wedges from the full optimal mechanism depicted in Figure 2. Panel C extends the age-dependent policy to four age groups. The pattern persists on average, but not exactly monotonically. Adding more detailed age brackets does not seem to improve the gain from the policy at all. The revenue loss is still the same and very large at 63-64%.

Overall, it appears that the strong nonlinearities highlighted in Section 5 are important components of the optimal mechanism and that shortcutting them costs a lot in terms of lost revenue.

Price subsidies on intermediate goods. We also consider the case of a pure price subsidy on monopolists' sales that aligns their private benefit from producing more quantity with the social one. The loss here relative to the full optimum is 66.4%, which is not much more than with the linear subsidy and tax added. Hence, the correction of the monopoly distortion is first order.

³⁸With linear instruments, all that is required is that the government can see some measure of aggregate quality or profits and keeps a constant share of it (equal to the linear tax rate). There is no need to observe the identity of individual firms and individual quantities.

TABLE 5: OPTIMAL LINEAR POLICIES

Age	Optimal τ	Optimal \tilde{s}	Optimal s	Revenue Loss
<i>A. Age-Independent Policies</i>				
1-15	31.4%	40.5 %	9.1%	65%
<i>B. Age-Dependent, 2 age brackets</i>				
1-15	31.6%	40.6%	9%	64%
16-30	31.0%	39.7%	8.7%	
<i>C. Age-Dependent, 4 age brackets</i>				
1-7	32.7%	40.8%	8.1%	63%
8-15	30.8%	39.8%	9%	
16-22	30.9%	40.1%	9.2%	
23-30	31.5%	40.15%	8.3 %	

Note: All optimal linear policies are computed assuming that the optimal price subsidy is in place, i.e., the linear profit tax is on profits post-price subsidy. Profits are taxed at rate τ gross of R&D expenses, i.e., R&D expenses are not deducted. Hence a firm's payoff is $\pi(1 - \tau) - (1 - \tilde{s})M(r)$. The equivalent net subsidy if R&D expenses were deductible from taxable corporate income is $s = \tilde{s} - \tau$.

Nonlinear, size-dependent R&D subsidy. Finally, we turn to a nonlinear subsidy rate, such that the subsidy received as a fraction of costs $M = M_t(r_t)$ is equal to $s(M)$ where

$$s(M) = c_0 + (c_1 - c_0) \cdot (1 - e^{-c_2 M})$$

At the optimum we find that the parameters should be set to $c_0 = 0$, $c_1 = 43\%$ and $c_2 = 46\%$. The optimal linear tax is set at $\tau = 31.56\%$. The welfare loss is now somewhat improved to 62% relative to the full unrestricted optimum. The optimal nonlinear subsidy is depicted in Figure 16 and is increasing and concave in R&D investment. Firms that spend more on R&D investments face a lower unit price for R&D investments, i.e., a higher marginal subsidy on their R&D.³⁹

³⁹This is in line with the idea of a "quantity discount" for high valuation consumers in the nonlinear monopolist pricing literature.

7.2 Comparative Statics

Table 6 recomputes the optimal linear tax and subsidy rates for different parameter values. As for the unrestricted policies in Section 6.4, we consider the three key parameters: the persistence p , the complementarity $\rho_{\theta r}$ and the strength of the technology spillover ζ . It is also clear that the gain from being able to use only restricted (here, linear) instruments strongly depends on parameters that affect the screening term.

A higher persistence slightly decreases the optimal linear tax and subsidy rates, but strongly reduces the (still large) revenue loss relative to the full optimum from using linear instruments. This is intuitive: the nonlinear, history-dependent full mechanism allows the planner to tailor the policies to firms' heterogeneous histories of productivity types. When productivity is more persistent, a firm of one type is more likely to remain that same type throughout time. The fine-tuning of the policies for different sized-firms is then much less valuable and linear tools are better at reaping the same benefits.

A higher complementarity of R&D investments with firm productivity increases the tax on profits and hence also decreases the net incentive on R&D investments, i.e., the effective subsidy rate post-deduction of R&D expenses equal to approximately $s - \tau = 5.7\%$. As explained theoretically in Section 5, a higher complementarity of R&D expenses with firm type means that firms' incentive constraints are tightened when R&D expenses are increased, which is costly. Here, with linear tax tools, there are no explicit incentive constraints, but the behavioral response of the firm to the taxes and subsidies follow the same logic.

Finally, a higher spillover strength decreases the optimal linear profit tax and increases the net subsidy on R&D from $4.9\% = 41.4\% - 36.5\%$ to $11.7\% = 39.7\% - 28\%$ so as to provide a higher Pigouvian correction.

Note an interesting pattern here related to the revenue losses from restricted instruments. The loss from linear policies is barely different when spillovers are large or small. This is because it is the screening term that one needs to tailor to the firm type – as in the unrestricted mechanism. As long as the Pigouvian correction – unrelated to the screening term – is on average at the right level given on the strength of the spillover, the revenue loss is the same. On the other hand, the other two parameters (persistence and complementarity) directly affect the screening term and, depending on their values, the revenue losses from the linear mechanism are very different.

8 Conclusion

In this project, we study how to most efficiently use tax policy to foster innovation. R&D tax policies in their current form are widely used, disparate, and costly, and have sometimes been criticized for a lack of effectiveness (Goolsbee, 2003). The market failures that warrant some policy intervention in the innovation process are the lack of appropriability of innovations and non-internalized technology spillovers between firms. Our core contribution is to consider asymmetric information. Firms'

TABLE 6: COMPARATIVE STATICS FOR THE OPTIMAL LINEAR POLICIES

Parameter	Optimal τ	Optimal \tilde{s}	Optimal s	Revenue Loss
<i>A. Role of Persistence p</i>				
$p = 0.5$	31.4%	40.5 %	9.1%	70.5%
$p = 0.9$	31.3%	40.3 %	9%	45.8%
<i>B. Role of Complementarity $\rho_{\theta r}$</i>				
$\rho_{\theta r} = \rho_{lr} = 1$	33.6%	41.3%	7.7%	61.7%
$\rho_{\theta r} = 1.2 > \rho_{lr}$	35.6%	41.3%	5.7%	58.9
<i>C. Role of the Technology Spillover</i>				
$\zeta = 0.01$	36.5%	41.4%	4.9%	73.8%
$\zeta = 0.03$	28%	39.7%	11.7%	73%

Note: All optimal linear policies are computed assuming that the optimal price subsidy is in place, i.e., the linear profit tax is on profits post-price subsidy. Profits are taxed at rate τ gross of R&D expenses, i.e., R&D expenses are not deducted. Hence a firm’s payoff is $\pi(1 - \tau) - (1 - \tilde{s})M(r)$. The equivalent net subsidy if R&D expenses were deductible from taxable corporate income is $s = \tilde{s} - \tau$.

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 new 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 micro-level firm data from COMPUSTAT matched to patent data to discipline and quantify our model.

Our main results are, first, to 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. 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

productivity: the more complementary R&D investment is to firm 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 productivity shocks and the strength of spillovers are other key determinants of the optimal policies. Second, we show that these constrained efficient allocations can be implemented with a parsimonious corporate income tax function and a price subsidy, or, equivalently, by a prize mechanism which depends on the observable innovation inputs and the innovation quality. Third, the estimation of our model based on key moments in the data allows us to quantify the optimal policies, as well as to show that the revenue losses from simpler policies, such as linear R&D and profit taxes, are very large, even if they are made conditional on firm age. This suggests that current policies could be much improved. Possible margins of improvement would be to condition R&D policies on some measure of innovation performance and by allowing for nonlinear policies.

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 innovation 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.


References

- ACEMOGLU, D., U. AKCIGIT, N. BLOOM, AND W. R. KERR (2013): "Innovation, Reallocation and Growth," Working Paper 18993, NBER.
- AGHION, P. AND P. HOWITT (1992): "A Model of Growth Through Creative Destruction," *Econometrica*, 60, 323–351.
- AKCIGIT, U., D. HANLEY, AND N. SERRANO-VELARDE (2014): "Back to Basics: Basic Research Spillovers, Innovation Policy and Growth," Working Paper 19473, NBER.
- AKCIGIT, U. AND W. R. KERR (2010): "Growth through Heterogeneous Innovations," Working Paper 16443, NBER.
- ALES, L. AND C. SLEET (2016): "Taxing Top CEO Incomes," *The American Economic Review*, 106, 3311–3366.
- ATKESON, A. AND A. BURSTEIN (2015): "Aggregate Implications of Innovation Policy," Working paper, UCLA.
- BARON, D. P. AND R. B. MYERSON (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50, 911–930.
- BLOOM, N., L. CHENNELLS, R. GRIFFITH, AND J. VAN REENEN (2002a): "How Has Tax Affected the Changing Cost of R&D? Evidence from Eight Countries," in *The Regulation of Science and Technology*, Springer, 136–160.
- BLOOM, N. AND R. GRIFFITH (2001): "The Internationalisation of UK R&D," *Fiscal Studies*, 22, 337–355.
- BLOOM, N., R. GRIFFITH, AND J. VAN REENEN (2002b): "Do R&D tax credits work? Evidence from a panel of countries 1979-1997," *Journal of Public Economics*, 85, 1–31.
- BLOOM, N., E. A. MAHAJAN, D. MCKENZIE, AND J. ROBERTS (2013a): "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, 128, 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 (2013b): "Identifying technology spillovers and product market rivalry," *Econometrica*, 81, 1347–1393.
- BLOOM, N. AND J. VAN REENEN (2007): "Measuring and Explaining Management Practices Across Firms and Countries," *Quarterly Journal of Economics*.
- 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.
- CHETTY, R. (2012): "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply," *Econometrica*, 80, 969–1018.
- CORNELLI, F. AND M. SCHANKERMAN (1999): "Optimal Patent Renewals," *RAND Journal of Economics*, 30, 197–213.
- CREMER, J. AND S. SCOTCHMER (1996): "Optimal R&D Procurement: Why Patents?" Working paper, University of California-Berkeley.
- FARHI, E. AND I. WERNING (2013): "Insurance and taxation over the life cycle," *The Review of Economic Studies*, 80, 596–635.

- GOLOSOV, M., A. TSYVINSKI, AND I. WERNING (2006): “New dynamic public finance: a user’s guide,” in *NBER Macroeconomics Annual 2006, Volume 21*, MIT Press, 317–388.
- GOLOSOV, M., A. TSYVINSKI, AND N. WERQUIN (2014): “A Variational Approach to the Analysis of Tax Systems,” Working Paper 20780, NBER.
- GOOLSBEE, A. (2003): “Investment Tax Subsidies and the Wages of Capital Goods Workers: To the Workers Go the Spoils?” *National Tax Journal*, LVI, 153–166.
- GRIFFITH, R., S. REDDING, AND J. REENEN (2001): “Measuring the Cost-Effectiveness of an R&D Tax Credit for the UK,” *Fiscal Studies*, 22, 375–399.
- GRILICHES, Z. (1998): *R&D and Productivity: The Econometric Evidence*, University of Chicago Press.
- 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. AND J. VAN REENEN (2000): “How Effective are Fiscal Incentives for R&D? A Review of the Evidence,” *Research Policy*, 29, 449–469.
- HICKS, J. (1970): “Elasticity of substitution again: substitutes and complements,” *Oxford Economic Papers*, 25, 289–296.
- JAFFE, A. AND M. TRAJTENBERG (2002): *Citations and Innovations: A Window on the Knowledge*, MIT Press.
- KLENOW, P. AND A. RODRIGUEZ-CLARE (2005): “Externalities and Growth,” in *Handbook of Economic Growth*, ed. by P. Aghion and S. Durlauf, Elsevier, 817–861.
- KREMER, M. (1998): “A Mechanism for Encouraging Innovation,” *Quarterly Journal of Economics*, 113, 1137–1168.
- LAFFONT, J.-J. AND J. TIROLE (1986): “Using Cost Observation to Regulate Firms,” *Journal of Political Economy*, 94, 614–641.
- LEAHY, D. AND J. P. NEARY (1997): “Public Policy Towards R&D in Oligopolistic Industries,” *The American Economic Review*, 87, 642–662.
- PAKES, A. (1985): “On Patents, R&D, and the Stock Market Rate of Return,” *Journal of Political Economy*, 93, 390–409.
- (1986): “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica*, 54, 755–784.
- PAKES, A. AND Z. GRILICHES (1984): “Patents and R and D at the Firm Level: A First Look,” in *R and D Patents and Productivity*, ed. by Z. Griliches, University of Chicago Press, 55–72.
- PAKES, A. AND M. SCHANKERMAN (1986): “Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period,” *Economic Journal*, 96, 1052–1076.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic mechanism design: A myersonian approach,” *Econometrica*, 82, 601–653.
- SAPPINGTON, D. (1982): “Optimal Regulation of Research and Development under Imperfect Information,” *The Bell Journal of Economics*, 13, 354–368.

- SCOTCHMER, S. (1999): “On the Optimality of The Patent Renewal System,” *RAND Journal of Economics*, 30, 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. (2014): “Learning and (or) Doing: Human Capital Investments and Optimal Taxation,” Working paper, Harvard University.
- (2016): “Optimal Taxation and Human Capital Policies over the Life Cycle,” *Journal of Political Economy*, forthcoming.
- 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,” Tech. rep.
- WERQUIN, N. (2016): “Income Taxation with Frictional Labor Supply,” Working paper, Toulouse School of Economics.


FIGURE 2: OPTIMAL PROFIT AND R&D WEDGES FOR DIFFERENT FIRM AGES




Notes: Panel (a) plots the average profit wedge $\bar{\tau}(\theta^t)$ at different ages of the firms. Panel (b) plots the average marginal R&D wedge $\bar{s}(\theta^t)$ at different ages. Panel (c) plots the average net marginal R&D wedge $s(\theta^t)$ at different ages.

FIGURE 3: OPTIMAL ALLOCATIONS AT DIFFERENT FIRM AGES

(a) R&D effort and investment




(b) Step size and profits




Notes: The left panel depicts the inputs, R&D expenses r_t and effort l_t , while the right panel depicts the outputs, the step size λ_t and the profits π_t , as a function of firm age on the horizontal axis.

FIGURE 4: OPTIMAL PROFIT AND R&D WEDGES FOR DIFFERENT FIRM TYPES




Notes: Panel (a) plots the optimal profit wedge $\tilde{\tau}(\theta^t)$ for $t = 2, 5, 10, 20$ for firms of different types. Panel (b) plots the optimal R&D wedge $\tilde{s}(\theta^t)$. Panel (c) plots the optimal net R&D wedge $s(\theta^t)$.

FIGURE 5: OPTIMAL ALLOCATIONS AT DIFFERENT FIRM QUALITIES




Notes: The left panel plots the optimal effort l_t for $t = 2, 3, 5, 10$ for firms of different types. The right panel plots the optimal R&D investment r_t .

FIGURE 6: OPTIMAL CORPORATE AND R&D WEDGES FOR DIFFERENT AGES WHEN QUANTITY CANNOT BE CONTROLLED




Notes: This figure illustrates the average profit wedge $\tilde{\tau}(\theta^t)$ (left panel) and the average R&D wedge $\tilde{s}(\theta^t)$ (right panel) at different ages when quantity cannot be controlled (i.e., intellectual property rights policy cannot be optimized on, patent protection is taken as given, and there is no price subsidy on the intermediate good).

FIGURE 7: OPTIMAL ALLOCATIONS AT DIFFERENT FIRM AGES WHEN QUANTITY CANNOT BE CONTROLLED




Notes: The left panel depicts the inputs, R&D expenses r_t and effort l_t , while the right panel depicts the outputs, the step size λ_t and the profits π_t , as a function of firm age on the horizontal axis when quantity cannot be controlled.

FIGURE 8: OPTIMAL CORPORATE AND R&D WEDGES FOR DIFFERENT FIRM QUALITIES WHEN QUANTITY CANNOT BE CONTROLLED




Notes: The left panel plots the optimal profit wedge $\tilde{\tau}(\theta^t)$ for ages $t = 2, 5, 10, 20$ for firms of different types. The right panel plots the optimal R&D wedge $\tilde{s}(\theta^t)$ when quantity cannot be controlled.

FIGURE 9: OPTIMAL ALLOCATIONS AT DIFFERENT FIRM QUALITIES WHEN QUANTITY CANNOT BE CONTROLLED




Notes: The left panel plots the optimal effort l_t for $t = 2, 5, 10, 20$ for firms of different types. The right panel plots the optimal R&D investment r_t when quantity cannot be controlled.

FIGURE 10: THE ROLE OF PERSISTENCE: WEDGES




Notes: The left panels plot the optimal profit wedge $\tilde{\tau}(\theta^t)$ for $t = 2, 5, 10, 20$ for firms of different types. The right panels plot the optimal R&D wedge $\tilde{s}(\theta^t)$. Panel (a) is for the case with low persistence $p = 0.5$, while panel (b) is for the case with high persistence $p = 0.9$.

FIGURE 11: THE ROLE OF PERSISTENCE: ALLOCATIONS




Notes: The left panels depict the R&D effort l_t for $t = 2, 5, 10, 20$ for firms of different types. The right panels depict the R&D investment r_t . Panel (a) is for the case with low persistence $p = 0.5$, while panel (b) is for the case with high persistence $p = 0.9$.

FIGURE 12: THE ROLE OF COMPLEMENTARITY: WEDGES




Notes: The left panels plot the optimal profit wedge $\tilde{\tau}(\theta^t)$ for $t = 2, 5, 10, 20$ for firms of different types. The right panels plot the optimal R&D wedge $\tilde{s}(\theta^t)$. Panel (a) is for the case $\rho_{\theta_r} = \rho_{r_l} = 1$, while panel (b) is for the case $\rho_{\theta_r} = 1.2 > \rho_{r_l} = 1$.

FIGURE 13: THE ROLE OF COMPLEMENTARITY: ALLOCATIONS




Notes: The left panels depict the R&D effort l_t for $t = 2, 5, 10, 20$ for firms of different types. The right panels depict the R&D investment r_t . Panel (a) is for the case $\rho_{\theta r} = \rho_{rl} = 1$, while panel (b) is for the case $\rho_{\theta r} = 1.2 > \rho_{rl} = 1$.

FIGURE 14: THE ROLE OF SPILLOVERS: WEDGES




Notes: The left panels plot the optimal profit wedge $\tilde{\tau}(\theta^t)$ for $t = 2, 5, 10, 20$ for firms of different types. The right panels plot the optimal R&D wedge $\tilde{s}(\theta^t)$. Panel (a) is for the case of a small spillover $\zeta = 0.01$, while panel (b) is for the case of a larger spillover $\zeta = 0.03$.

FIGURE 15: THE ROLE OF SPILLOVERS: ALLOCATIONS



Notes: The left panels depict the R&D effort l_t for $t = 2, 5, 10, 20$ for firms of different types. The right panels depict the R&D investment r_t . Panel (a) is for the case of a small spillover $\zeta = 0.01$, while panel (b) is for the case of a larger spillover $\zeta = 0.03$.

FIGURE 16: OPTIMAL NONLINEAR R&D SUBSIDY



Notes: The figure depicts the optimal nonlinear subsidy rate of the form $s(M) = c_0 + (c_1 - c_0) \cdot (1 - e^{-c_2 M})$ such that the agent only pays cost $(1 - s(M(r))) \cdot M(r)$ for an investment $M(r)$. At the optimum we find that the parameters should be set so: $c_0 = 0$, $c_1 = 43\%$ and $c_2 = 46\%$.

Appendix

A.1 Additional Results

A.1 Simple Model

Tax implementation: The first-order conditions of the firm with the 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:

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

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

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: We can similarly derive the relation between the marginal grant and the wedges, under the assumption that the grant function is differentiable in its two arguments.

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

A.2 Functional form example with constant markups

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:

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

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

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

A.2 Proofs of the Propositions in the Main Text

Proof of Propositions 1 and 2:

Taking the first-order conditions of the planner's problem in (7) 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 4:

Taking the FOC of program P in (17) with respect to $r_t(\theta^t)$ yields:

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

Using the definition of the R&D wedge as:

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

to substitute for the marginal cost $M'_t(r_t(\theta^t))$ in the FOC, we obtain formula (22).

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

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

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 Proposition 5:

For every period, define the following objects:

$$D_s(\theta^{s-1}, \theta_s) = E \left(\sum_{t=s}^T I_{(s),t} \frac{\partial \tilde{U}}{\partial \theta_t} | \theta^t \right) \\ Q_s(\theta^{s-1}, \theta_s) = \int_{\underline{\theta}}^{\theta_s} D_s(\theta^{s-1}, q) dq$$

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

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

In the unrestricted mechanism, the transfers provided every period are:

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

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 manager 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 equilibrium outcome (in the status quo case or for various policy experiments), 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} , one must employ both forward and backward iteration.

Given a certain candidate path for \bar{q} , we can find the optimal choices for research spending and manager 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, it is useful to dampen the updating process to avoid any instabilities.

To generate simulated moments for parameter estimation, we simulate a large number of firms 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 spillover regression coefficient, which is used to identify the externality parameter.

For that moment, we actually re-solve and re-simulate the model for a variety of different values for κ_r (the research cost scale parameter) clustered around the true κ_r value. We interpret each simulated economy as representing a particular industry with a particular value for κ_r , or, alternatively, with a different R&D subsidy or tax credits. This mimics the exogenous variation used to identify the spillovers in the Bloom et al. (2013b) paper. Using this variation in cost, we then run a regression of firm sales on the cumulative amount of research spending the firm has done (a knowledge stock of sorts) as well as the average research spending by all firms in that period and industry. We then match this to an analogous run by Bloom et al. (2013b).

A.2 Identification and the Jacobian Matrix

TABLE A1: JACOBIAN MATRIX FOR PARAMETER ESTIMATION

	ρ	α	κ_r	κ_l	p	σ_ε	Θ_1 range	ζ
M1. Patent quality-R&D elasticity	0.01	0.00	0.00	0.00	0.02	0.01	0.01	0.00
M2. R&D/Sales median	0.00	0.05	0.00	0.00	0.00	0.00	-0.02	0.00
M3. Sales growth	0.01	0.01	-0.01	-0.02	0.01	0.01	-0.01	0.01
M4. Within-firm patent quality coeff of var	0.00	-0.01	0.00	0.00	0.02	0.02	0.02	0.00
M5. Across-firm (young) coeff of var	0.00	-0.02	0.00	0.00	0.06	0.01	0.04	0.00
M6. Across-firm (old) coeff of var	0.00	-0.02	0.00	0.00	0.09	0.05	0.01	0.00
M7. Patent quality young/old	-0.02	-0.08	0.03	0.04	-0.05	-0.02	0.07	-0.01
M8. Spillover coefficient	0.13	0.08	-0.01	-0.02	-0.06	-0.03	0.01	0.01

Note: This table reports the percentage change in the moment (row) for a 5% change in the parameter (column) from its baseline value, while keeping the rest of the parameters at their benchmark values. We report the average of the +5% and -5% changes. Therefore, these values can be interpreted as a double-sided discrete approximation.