Subsidizing Research Programs with "If" and "When" Uncertainty in the Face of Severe Informational Constraints*

David Besanko[†] Jian Tong[‡] Jason Jianjun Wu[§]

Northwestern University University of Southampton Compass Lexecon

October 11, 2017

Abstract

We study socially optimal subsidy policies for research programs when firms have private information about the likelihood of project viability, but the government cannot form a unique prior belief about this likelihood. If the shadow cost of public funds is zero, the government can attain the first-best level of welfare as a (belief-free) *ex post* equilibrium through simple policies: a pure matching subsidy under monopoly R&D and a matching subsidy combined with an unrestricted grant under R&D competition. If the shadow cost of public funds is positive, first-best welfare cannot be attained as an *ex post* equilibrium under either monopoly or competition. However, max-min subsidy policies exist under both scenarios and consist entirely of pure matching subsidies. Allowing firms to form an R&D consortium reduces the matching rate for the highest max-min subsidy relative to R&D competition, while increasing both R&D investment intensity and cumulative investment. Enabling cooperative R&D can both economize on the shadow costs of public subsidies and strengthen investment incentives.

Keywords: Research and development, subsidies, dynamic stochastic games, asymmetric information, belief-free games, social cost of public funds, research consortia.

JEL classification: O38, D60, D82, H2, C73.

1 Introduction

Governments have long played a role in subsidizing private investments in R&D. The principal economic justification for R&D subsidies is the presence of market imperfections (e.g., limits on appropriability

*The authors would like to thank Alberto Galasso and Nick Klein for their very helpful comments as well as participants at 2009 International Industrial Organization Conference at Boston and 2010 Southwest Economic Theory Conference at Los Angeles. We also thank Mark Armstrong and two anonymous referees for their valuable suggestions. Besanko gratefully acknoweldges the financial support from the National Science Foundation under Grant 0615615.

[†]Department of Strategy, Kellogg School of Management, Northwestern University, Evanston, IL 60208. Email: d-besanko@kellogg.northwestern.edu

[‡]Economics Department, School of Social Sciences, University of Southampton, Southampton, SO17 1BJ, UK, Email: j.tong@soton.ac.uk

[§]Compass Lexecon, 105 College Road East, Princeton, NJ 08540. Email: jwu@compasslexecon.com.

or problems of free riding) that result in socially suboptimal provision of private R&D (Arrow, 1962). But the economic literature presents mixed results on how effectively subsidies can address these market failures. Some empirical studies conclude that R&D subsidies stimulate private R&D investment (e.g., Lach, 2002 and Almus and Czarnitzki, 2003), while others find that subsidies crowd out private R&D investment (e.g., Irwin and Klenow, 1996 and Wallsten, 2000) or leave it unchanged (e.g., Klette and Moen, 1998, 1999). The theoretical literature (discussed in more detail below) shows that subsidies can, in principle, stimulate private R&D investment and increase social welfare, but the models in the literature are typically cast within static and/or deterministic settings that seem far removed from the dynamic and uncertain environments in which much modern research (especially basic research) takes place. From both a theoretical and empirical perspective, the effectiveness of government R&D subsidies remains an unsettled question, and as Hall (2005) suggests, it deserves further research.

The purpose of this paper is to advance the theory of R&D subsidies by studying their impact in a setting with two key features: (1) firms undertake a research program that involves uncertainty about the *timing* of the scientific breakthrough that the program can lead to ("when uncertainty") and about the underlying *viability* of the program ("if uncertainty"); (2) firms' prior beliefs about program viability are private information, and the government providing the subsidies is unable to form a unique prior belief about the firms' priors. That is, the subsidizing government faces a severe informational asymmetry—it neither knows how optimistic firms are about the viability of the research program, nor does it know what to believe about firms' optimism. These features seem especially likely to hold for research programs taking place in "uncharted waters," so our theory is particularly relevant for groundbreaking research programs in areas in which there is little established consensus about whether the current direction of inquiry is likely to be fruitful.

Our paper makes two contributions. First, we show that the *way* in which R&D is subsidized matters. Certain types of subsidies (i.e., earmarked funding in which a firm is required to spend a certain minimum amount on the project) may suppress private investment, while other types of subsidies (i.e., a pure matching subsidy in which a firm undertaking R&D is reimbursed a fraction of its R&D expenses) may stimulate private investment. This suggests that empirical studies of how subsidies impact private R&D investment need to be cognizant of how subsidies are structured. Second, we show that despite the government's severe informational constraints, a simple mechanism—in particular a pure matching subsidy—can still perform reasonably well with respect to plausible decision criteria, and under interesting circumstances it can even attain the first-best level of welfare.

More specifically, our paper uses a two-armed bandit model of R&D investment in which firms seek to achieve a scientific breakthrough.[1] As time passes and the breakthrough is not achieved, the firms' posterior probability that the project is viable falls. In the absence of a subsidy, once that posterior reaches an abandonment threshold—which depends on the payoff from achieving the breakthrough, the marginal cost of R&D investment, and the rate at which investment increases the likelihood of a viable project achieving a breakthrough—firms stop investing and terminate the project. To incentivize private R&D investment, the government provides firms with an R&D subsidy that depends on their R&D investments. The subsidy mechanism that we consider subsumes three specific funding schemes commonly used in practice: a pure matching subsidy, an earmarked subsidy, and an unrestricted subsidy in which the government makes an open-ended commitment to fund the project until a breakthrough occurs (though unlike an earmarked subsidy there is no formal requirement that the firm actually spend the money on the focal R&D project).

We consider two cases: monopoly R&D, in which the breakthrough is pursued by a single firm, and R&D competition, in which $N > 1$ firms compete to achieve the breakthrough. Under monopoly R&D, the matching component of the subsidy reduces the abandonment threshold, thus expanding the range of posterior beliefs over which the monopolist undertakes the maximum feasible R&D investment, while the earmarked and unrestricted components of the subsidy increase the abandonment threshold, shrinking the range of maximum investment. Under R&D competition, a complication arises that does not exist under monopoly: the possibility that firms may free ride on the R&D efforts of other firms. In this setting, the unrestricted component to the subsidy and the minimum mandated level of R&D—which had unambiguously adverse incentive effects under monopoly—can eliminate free riding. Indeed, without unrestricted funding or a minimum mandate, the free-rider problem always arises (in the absence of an explicit or tacit research consortium).

We then turn to optimal subsidy policy in a setting in which firms have private information about the prior likelihood that the project is viable, and the government lacks probabilistic knowledge of this likelihood; in other words, the government has no (unambiguous) prior over the firm's prior. This

[1] The use of two-armed bandit models to analyze economic problems dates back to Rothschild (1974). Recently, a number of papers focus on the strategic interaction among agents in a bandit framework (e.g., Keller, Rady, and Cripps, 2005, and Klein and Rady 2011). Our paper is closest to Keller, Rady, and Cripps (2005), as our second stage R&D competition is based on their Poisson bandit framework. Besanko and Wu (2013) explore R&D competition and cooperation in a model inspired by Keller, Rady, and Cripps (2005).

severely constrained informational environment prevents the government from using either a forcing mechanism to achieve first-best welfare or a conventional mechanism design approach in which it offers the firms a menu of policies that maximizes expected social welfare for a given prior probabilistic belief over the firms' private information. When the shadow cost of public funds is zero, under monopoly R&D there exists a pure matching subsidy that induces the firm to follow the first-best R&D policy irrespective of its prior beliefs about the viability of the project. This particular matching subsidy is thus a (belief-free) *ex post* equilibrium. The first-best outcome can also be achieved under R&D competition. However, unlike monopoly, that policy is not a pure matching policy. Instead, it involves a combination of a matching subsidy and an unrestricted subsidy. The unrestricted component of the subsidy eliminates the free-rider problem, and given this, the matching rate is set to mimic the social planner's optimal investment policy.

By contrast, when there is a positive shadow cost of public funds, we prove that an *ex post* equilibrium does not, in general, exist either under monopoly R&D or R&D competition. We then consider a policy-making objective that is appropriate for belief-free games:[2] the max-min criterion (Gilboa and Schmeidler, 1989). Max-min policies protect the policy maker against inferior "worst case scenario" policy outcome relative to any alternative policies. This criterion strikes us as plausible inasmuch as policy makers operating under significant informational constraints may very well be keen to avoid big policy "mistakes." We show that under both monopoly R&D and R&D competition, the set of (pure strategy) max-min policies are pure matching subsidies. Under monopoly R&D, the policy in this set with the highest matching rate has the appealing property of eliminating the possibilities of both underinvestment and overinvestment in R&D. Further, as the shadow cost of public funds goes to 0, this matching rate approaches the matching rate that implements the first-best solution. In contrast to monopoly R&D, the policy under R&D competition in the max-min set with the highest matching rate cannot overcome suboptimal intensity of investment due to the free-rider problem. This suggests that the government might benefit by enabling the firms to choose their R&D investments cooperatively. We show that if the government permits firms to choose R&D cooperatively through a research consortium, there exists a matching rate that satisfies the max-min criterion and eliminates both over- and underinvestment. This matching rate is less than the highest max-min matching rate under R&D competition, and it induces a lower abandonment threshold than the one arising under

[2] See Bergemann and Morris (2007) for a general treatment of belief-free incomplete information games.

R&D competition when firms receive the highest max-min subsidy rate. Thus, an R&D consortium can stimulate R&D investment, while economizing on the shadow cost of public funds.

Our paper fits within the theoretical literature on R&D subsidy policy. Papers in this literature have focused on a number of broad issues. Some, such as Spencer and Brander (1983) and Qiu and Tao (1998), study the use of R&D subsidies to enhance national competitiveness. Other papers consider the role of subsidies to help overcome informational problems. For example, Socorro (2007) explores optimal patent subsidies for R&D in the context of a mechanism design problem in which firms have private information about the value of an uncertain R&D project, while Takalo and Tanayama (2010) examine whether R&D subsidies can alleviate financing constraints due to adverse selection. Most closely related to this paper are papers by Hinloopen (1997, 2000), Stenbacka and Tomback (1998), and Romano (1989) that explore the impact of subsidies on the level of R&D and social welfare. Hinloopen (1997) analyzes a model similar to the framework of d'Aspremont and Jacquemin (1988) and Kamien et. al. (1992) in which firms' investments in cost-reducing effort deterministically reduces their costs, and possibly the costs of other firms as well due to spillovers. R&D subsidies are shown to increase the level of investment activity and social welfare and are more effective at increasing R&D investment than allowing firms to cooperate through research joint ventures or R&D cartels. Stenbacka and Tomback (1998) analyze the best way to organize R&D (e.g., competitive research joint venture, cartelized research joint venture, R&D competition) given that the government chooses an optimal subsidy rate for the mode of organization being considered. With optimal subsidies, research joint ventures are shown to be superior to competition provided that the social cost of subsidies is not too large. Romano (1989) analyzes subsidies for research projects aimed at achieving process innovations in the presence of "when" uncertainty. He shows that it is always socially optimal to subsidize a monopolist, but under certain circumstances (e.g., sufficiently long patent life) it is not optimal to subsidize competitive firms.

Our paper differs from the existing theoretical literature in several important respects. First, unlike the existing literature that analyze R&D subsidies in reduced-form static or two-stage models, our model is explicitly dynamic. By employing the two-armed bandit framework, we can analyze how alternative subsidy policies affect belief updating and the abandonment of R&D projects, issues that cannot be studied in static or two-stage models. Second, we consider more general subsidy policies than those considered in the existing literature. Hinloopen (1997, 2000) and Stenbacka and Tomback (1998) consider pure matching subsidies, while Romano considers unrestricted subsidies. Our paper, by contrast, analyzes a more general subsidy mechanism that embraces both matching and unrestricted subsidies as special cases. Third, in contrast to many of the papers cited above, a key focus of our paper is on the properties of an optimal subsidy policy and how that policy is affected by underlying economic fundamentals. Finally, we consider an environment in which the policy maker lacks prior beliefs over firms' private information about project viability. Accordingly, optimal subsidy policy in our model cannot depend on the details of potentially *ad hoc* subjective beliefs and must instead be robust to the entire range of possible assessments that the firms might have about project viability.

Our paper is also related to several papers in the broader literature on the financing of innovative activity, in particular Bergemann and Hege (1998, 2005) and Hörner and Samuelson (2013). These papers, like ours, study R&D projects that are characterized by both "if" and "when" uncertainty. The main focus of these papers is to explore the hidden action (and its induced hidden information) problems in a context of open-ended exclusively external financing. Although it is critical to understand the problem caused by hidden action in R&D experimentation, it is equally important to understand the economics of R&D subsidies in the case of severe information asymmetries. By shifting the emphasis from hidden action to the hidden information problem in the context of a belief-free incomplete information game, our paper complements the current literature on funding experimentation by including a number of new important features. First, unlike exclusively external funding, we allow firms to use their own funding to pursue R&D while receiving financial support from the government. Second, we address the appropriability and free-riding problems simultaneously with a funding scheme that includes matching, earmarked, and unrestricted subsidies as special cases, while the current literature largely restricts the external financial support to unrestricted funding only. Third, we study the funding policy for multiple firms while the Bergemann and Hege and Hörner and Samuelson papers focus on the one-firm case. Finally, their models assume no friction in external funding, while we consider cases with and without frictions in terms of a shadow cost.

Though our paper is the first to study optimal research subsidization under ambiguity, it can also be connected to literatures on public policy or mechanism design that impose informational restrictions reflecting ambiguity or ignorance. For example, by way of acknowledging model misspecification in macroeconomic theory, Hansen and Sargent (2001a, 2001b) introduce a robust control approach which is equivalent to entertaining a set of multiple priors and adopting the max-min expected utility as an objective function akin to Gilboa and Schmeidler (1989). An implication of this work is that a “preference for robustness induces context-specific precaution. In asset pricing models, this boosts market prices of risk. … In permanent income models, it induces precautionary savings. In sticky-price models of monetary policy, it can induce a policy authority to be more aggressive in response to shocks.” Precaution induced by a preference for robustness can be found in other applications, such as climate change policy (Xepapadeas, 2012) and auction design (Lo, 1998, Bose, Ozdenoren, Pape, 2006), as well as mechanism design in general (Bodoh-Creed, 2012, De Castro and Yannelis, 2010). For example, using numerical examples, Xepapadeas illustrates that as the level of ambiguity (i.e., model misspecification error) increases, the optimal precautionary reduction of emission also increases. Precaution can be seen at work in our model. When the shadow cost of funds is positive, the pure matching subsidies that constitute the max-min set compensate firms only for the R&D effort they undertake. This avoids the costly possibility that firms are excessively subsidized for an R&D project with a low likelihood of viability, which could happen under an unrestricted or earmarked subsidy.

The paper is divided into five sections, including this one. Section 2 describes the model. Section 3 illustrates the incentive effects of subsidy policies. Section 4 takes up the question of how the government, lacking both deterministic and probabilistic knowledge of firms’ priors about the viability of the project, would determine an optimal R&D subsidy policy. Again, we start with the case of monopoly R&D and then turn to R&D competition. We also analyze the potential social benefits from formation of a research consortium. Section 5 summarizes and concludes. Proofs of all propositions are in the Appendix.[3]

2 The Model

We present a model of R&D investment based on the exponential bandit framework of Keller, Rady, and Cripps (2005). We state the model with N firms, with the analysis of monopoly corresponding to the special case of $N = 1$.

Each of the N identical firms faces an opportunity to invest in an R&D program aimed at achieving a significant breakthrough. *Ex ante* the firms do not know if a breakthrough is possible. Let p_0 denote the firms’ common prior that the project is viable, i.e., that the breakthrough can be achieved eventually. Conditional on the project being viable, the time the breakthrough occurs is random. Higher R&D investment increases the likelihood that the breakthrough occurs sooner. Specifically, let

[3] An Online Appendix contains detailed derivations of the equilibria under monopoly R&D and R&D competition, as well as further discussion of the max-min decision criterion.

k_t^i, $i = 1, \ldots N$, denote firm i's R&D investment at time t. Conditional on the project being viable, the hazard rate of firm i's success on the R&D project is λk_t^i, where $\lambda > 0$ is a parameter. We assume that each firm faces a technological constraint that limits its investment in R&D to at most 1 unit of effort at any point in time t, meaning $k_t^i \in [0, 1]$. If the project was indeed viable, and a single firm exerted the maximum feasible level of R&D effort, then $\frac{1}{\lambda}$ would be the expected time until a breakthrough occurs. R&D effort is costly, and the total cost $C(k_t^i)$ of R&D effort is assumed to be an identical linear function for each firm, $C\left(k_t^i\right) = \alpha k_t^i$, where $\alpha > 0$ denotes the marginal cost of R&D effort.[4]

A breakthrough is assumed to be "big news" and thus visible to all firms competing in the R&D race. As time passes and a breakthrough has not occurred, firms become more pessimistic about the viability of the project. Let $p(t)$ denote firms' posterior belief about the project's viability at date t. If no breakthrough occurs, $p(t)$ adjusts downward consistent with Bayes rule. The rate of belief updating is given by

$$dp = -\lambda \sum_{i=1}^{N} k_t^i p (1-p)\, dt. \tag{1}$$

The solution concept is Markov Perfect Equilibrium, with each firm's common posterior belief p being the payoff-relevant state variable and equation (1) representing the law of motion for the state variable.[5] Investment behavior and firm value functions are thus conditioned on p. It is straightforward to establish that for the important case of constant "flat out" investment, i.e., $k_\tau^i = 1$ for $\tau \in [0, t]$ and all $i = 1, \ldots N$, the posterior belief about project viability if no breakthrough has occurred by date t is given by $p(t) = \frac{p_0}{p_0 + (1-p_0)e^{\lambda t}}$.[6] When the prior belief is close to 1, the posterior belief evolves very slowly for a long period of time if no breakthrough occurs, and in the extreme case of $p_0 = 1$—i.e., there is no "if" uncertainty—beliefs remain at $p_0 = 1$ even as time passes without a breakthrough.[7]

We assume the firm that wins the R&D race earns a payoff $\Pi > 0$. This payoff is the present value of the winning firm's profits, which are assumed to be discounted at a rate r. Each of the $N-1$ non-winning firms is assumed to receive a payoff $\theta\Pi$, where $\theta \in [0, 1]$. If $\theta = 0$, the R&D race is winner-take-all;

[4] The linearity of the cost function is needed to solve for the equilibrium investment level in closed form. The basic intuition underlying the results does not depend on the linearity of the cost function.

[5] Our use of the term Markov Perfect Equilibrium does not restrict it to be a refinement of subgame perfect Nash equilibrium. This relaxation allows us to apply it to our particular dynamic game with incomplete information. Since the government is the first mover who moves only once, to satisfy the requirement of Markov perfection, we only need each firm's strategy to be Markovian, and solve the appropriate dynamic programming problem.

[6] As we will see, constant investment, at least for a while, occurs along the equilibrium path for both $N = 1$ and $N > 1$.

[7] Suppose for example, if $p_0 = 0.9999$ and $\lambda = 0.1$. If firms are investing "flat out," then by $t = 50$, the posterior is still greater than 0.98.

if $\theta > 0$, the breakthrough has positive spillovers. The spillover parameter θ has a number of possible interpretations. It could reflect the degree to which the breakthrough is imitable by rivals in the absence of patent protection, or it could represent time discounting from the date at which patent protection from the breakthrough ends (so a smaller θ corresponds to a longer patent length). The discounted present value of the social benefit from the new technology is given by $CS + \Pi + (N-1)\theta\Pi$, where CS is the benefit of the breakthrough for consumers that the discovering and non-discovering firms cannot capture.[8] For later use, let $\rho \triangleq \frac{\Pi+(N-1)\theta\Pi}{CS+\Pi+(N-1)\theta\Pi} \in (0,1]$ be the appropriability ratio, i.e., the share of the total benefit captured by firms and thus $1-\rho$ is the share of the total benefit that accrues to consumers. Therefore, government intervention in the market for R&D has two potential justifications: the presence of R&D spillovers across firms (when $\theta > 0$) and the imperfect appropriability of the benefit from the breakthrough (when $\rho < 1$). Throughout the analysis, we assume that $\frac{\lambda[CS+\Pi+(N-1)\theta\Pi]}{\alpha} > 1$, which implies that the social benefit-cost ratio of a viable R&D project exceeds 1.[9]

The government's R&D subsidy policy is represented by an array of three instruments, $\mathbf{a} = (z, s, \phi)$, that form a schedule $S(k_t^i|\mathbf{a})$ that determines the funding flow a firm receives at each instant in time prior to a breakthrough:

$$S(k_t^i|\mathbf{a}) = \begin{cases} s + \phi\alpha(k_t^i - z) & \text{if } k_t^i \geq z, \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

In this schedule:

- $z \in [0,1]$ is the minimum R&D effort mandated by the government at each instant in time in order for the firm to be eligible for any funding, and thus αz is the minimum mandated spending on R&D.

- $s \in [\alpha z, \bar{s}]$ is the baseline amount of funding the firm receives, provided it satisfies the mandate, where $\bar{s}$ is a finite limit of s assumed to be such that $\frac{\bar{s}}{r} < \Pi$.[10] We require that $s \geq \alpha z$ so that

[8] If the N R&D competitors become N product market competitors, then Π would be a non-increasing function of N and CS would be a non-decreasing function of N. We suppress these dependences in much of what follows, but we return to them when we discuss comparative statics with respect to N.

[9] Let $\widetilde{T}$ be the random time to discovery for a project that is certain to be viable. With hazard rate λ and flat-out investments (by all N firms) at any point in time, $\widetilde{T}$ is an exponential random variable with parameter $N\lambda$. The *ex ante* expected social benefit of a viable R&D project would be $[CS + \Pi + (N-1)\theta\Pi]\, E(e^{-r\widetilde{T}})$, which equals $\frac{N\lambda[CS+\Pi+(N-1)\theta\Pi]}{r+N\lambda}$. The *ex ante* expected cost of a viable R&D project would be $\int_0^\infty N\alpha \left(\int_0^t e^{-r\tau} d\tau\right) \lambda e^{-N\lambda t} dt$ which can be shown to equal $\frac{N\alpha}{r+N\lambda}$. The *ex ante* benefit cost ratio is thus $\frac{\lambda[CS+\Pi+(N-1)\theta\Pi]}{\alpha}$.

[10] This implies that an unrestricted subsidy can never be so large that a firm would (weakly) prefer collecting the subsidy to receiving the prize with certainty.

at any point in time the firm would prefer to adhere to the mandate and accept the associated funding, rather than reject it. If $s = \alpha z$, the government exactly reimburses the firm its mandated R&D spending, while if $s > \alpha z$, the firm receives a subsidy in excess of its minimum mandated R&D expenditure. In this latter case, the firm could (in principle) spend some of its government funds on other activities besides the focal R&D project (e.g., it could fund other R&D projects). We thus refer to $s - \alpha z$ as the *unrestricted component* of its R&D subsidy.

- $\phi \in [0, 1]$ is the matching rate: the additional funding the firm receives for every additional dollar of R&D spending undertaken above the mandated level.

Throughout we let $A \triangleq \{\mathbf{a} | z \in [0, 1], s \in [\alpha z, \bar{s}], \phi \in [0, 1]\}$ denote the set of feasible subsidy policies. The set of policies in A embrace three interesting special cases:

- If $z = 0$, $s = 0$, and $\phi \in (0, 1]$, then a firm receives a *pure matching subsidy*: for every αk dollars of R&D investment, the government matches the firm's R&D spending by providing a subsidy of $\phi \alpha k$.

- If $z > 0$, $s = \alpha z$, and $\phi = 0$, then a firm receives an *earmarked subsidy*: it receives a subsidy of αz dollars, provided that its R&D effort satisfies the mandate z.

- If $z = 0$, $s > 0$, and $\phi = 0$, then a firm receives a *pure unrestricted subsidy*: it receives a no-strings-attached grant of s.

We assume throughout that the experience and scientific facts that underpin p_0 are unknown to the government and are thus private information to the firms. Therefore, p_0 can be interpreted as the firms' unobservable type. Lacking knowledge of p_0, the government cannot infer the firms' posterior belief $p(t)$ and, therefore, cannot write a "forcing contract" in which it ties the parameters of the subsidy function to the posterior in such a way that it replicates the first-best investment policy.[11] We further assume, in contrast to the conventional Bayesian approach in which the government would have a given prior belief about p_0, that the government has no such prior. In other words, the government not only lacks complete information about p_0, but it also lacks probabilistic information about it. Without a prior over p_0, the government cannot determine an optimal menu of policies that maximizes its expected

[11] The first-best investment policy is stated in Section 4.1.

welfare, as in a conventional mechanism design approach. An advantage of this assumption is that it forces consideration of robust policies that would not need to be fine-tuned in the wake of changes in potentially arbitrary prior beliefs.

Note that we focus exclusively on *ex ante* subsidies that are paid during the research phase of the project and do not consider *ex post* subsidies that are contingent on the success of the project (sometimes called patent subsidies; see Socorro, 2007). We justify this focus because *ex ante* subsidies are the most common way that governments support private R&D activity.[12]

We also restrict attention to stationary subsidy policies. A rationale for this restriction is as follows. Because the government faces a severe informational constraint—its priors about the firms' type are not informative—any learning about the unknown type p_0 would have a limited effect on altering the problem the government is solving.[13] Thus, from the government's point of view, the policy design problem can be approximated as a stationary one, which can be solved by a stationary policy scheme of the type we consider. This approximation is, to be sure, imperfect, but it has the advantage of technical tractability. When we restrict attention to stationary subsidy schemes, we can derive a closed-form solution of the monopoly firm's optimal response to the government's policy; we can partially characterize the symmetric equilibrium response to government policy under R&D competition when $N > 1$, and we can obtain clear analytical results on the optimal policy (within the restricted class). Further, as we show below, restricting attention to stationary policies may entail no loss of generality since, under important circumstances, they are powerful enough to attain the first-best outcome.[14]

[12] In our model, patent subsidies would enable the government to achieve the first-best solution under monoply when the shadow cost of public funds is zero. Even in this case, though, they would require that the government pass along the entire consumer surplus to the winning firm. This may be possible if the government itself is the only consumer of the products created by the research, as in the case of defense-related R&D. But this could be difficult in the context of other crucial technologies such as stem cell research.

[13] More specifically, from the observed effort of the firm, which may merely be a stream of binary choices, the government can only infer the infimum and/or the supremum of the type of the firm. This limited information has little value for conditioning the subsidization scheme. Furthermore, if the government has commitment to the given policy, i.e., it is negotiation-proof, then the learning about the firms' type has no value. If the government cannot commit to the policy, i.e., it is not negotiation-proof, then we will have to impose the condition of sequential rationality on the government policy, which makes the analysis intractable.

[14] A particular example of a time-varying subsidy policy would be one in which funding is cut off after a certain deadline. Bonatti and Hörner (2011) analyze the role of deadlines in collaborative research, and they show how deadlines can overcome the moral hazard in teams. In the concluding section, we discuss the role that deadlines might play in our model and how our model relates to the insights developed by Bonatti and Hörner.

3 The Incentive Effects of R&D Subsidies

To set the stage for the analysis of the optimal subsidy mechanism, in this section we discuss the incentive effects of R&D subsidy schemes.

3.1 Monopoly R&D ($N = 1$)

We begin by studying the case in which a single firm conducts the R&D project. The derivation of the optimal solution in this case parallels the characterization of the cooperative solution in Proposition 3.1 in Keller, Rady, and Cripps (2005), so we omit it here.[15] Given a subsidy policy $\mathbf{a}$, the firm's optimal R&D strategy can be characterized by an *abandonment threshold* $p_1(\mathbf{a})$ given by[16]

$$p_1(\mathbf{a}) = \frac{\alpha\,(1-\phi)}{\lambda\left[\Pi - \left(\frac{s-\alpha z}{r}\right)\right]\left(\frac{r}{r+\lambda z}\right)} > 0. \tag{3}$$

For any arbitrary belief $p \in [0,1]$, the monopolist's optimal R&D strategy is

$$k_1(p) = \begin{cases} 1 & \text{if } p > p_1(\mathbf{a}) \\ z & \text{if } p \leq p_1(\mathbf{a}) \end{cases}.$$

The monopolist's optimal investment decision $k_1(p)$ is a "bang-bang" rule: k either equals the minimum required level z, or the maximum feasible level 1, depending on whether its belief p about the project's viability is greater or less than the abandonment threshold.

A firm's equilibrium path behavior under the optimal policy depends on its prior belief about the project's viability, p_0. If $p_0 > p_1(\mathbf{a})$, i.e., the firm is sufficiently optimistic about the project's viability *ex ante*, there exists a non-empty interval $(p_1(\mathbf{a}), p_0]$ of posterior beliefs p over which it exerts maximum R&D effort (i.e., $k = 1$), an interval we refer to as the *range of maximum investment*.[17] In this case, the firm begins by investing flat out. As time passes without a breakthrough, it becomes more pessimistic about the viability of the project. Once its posterior falls to $p_1(\mathbf{a})$, the firm switches from "flat out" investment to the minimum level mandated by the government (i.e., $k = z$). If, on the other hand, $p_0 \leq p_1(\mathbf{a})$, i.e., the firm is sufficiently pessimistic about the viability of the project *ex ante*, the range

[15] For interested readers, a detailed derivation is avalable in the Online Appendix that accompanies this paper.
[16] Because $\frac{\bar{s}}{r} < \Pi$, it follows that $\Pi - \frac{s-\alpha z}{r} > 0$ for all feasible s and z, and thus $p_1(\mathbf{a}) > 0$.
[17] Clearly, this case cannot arise if $p_1(\mathbf{a}) \geq 1$.

of maximum investment is empty, and the firm begins by exerting the minimum mandated level z and continues doing so as long as a breakthrough does not occur.

The abandonment threshold is derived from a "marginal cost equals marginal benefit" condition. The marginal cost of an additional unit of R&D effort above the minimum threshold is $(1-\phi)\alpha$. The marginal benefit $\lambda p\left(\Pi - \frac{s-\alpha z}{r}\right)\frac{r}{r+\lambda z}$ consists of two components: the incremental increase in the likelihood of a breakthrough, λp, and the net prize to the firm if a breakthrough occurs, $\left(\Pi - \frac{s-\alpha z}{r}\right)\frac{r}{r+\lambda z}$. The firm's net prize is the present value of profits from a breakthrough, Π, minus the present value of the fungible portion of the subsidy, $\frac{s-\alpha z}{r}$, which the firm foregoes if it achieves a breakthrough. The net prize is further "deflated" by $\frac{r}{r+\lambda z}$, which equals the rate of percentage reduction in the expected present value of \$1 per one percent reduction in R&D effort, when that effort falls from $k=1$ to $k=z$.[18] As z increases, the extent of this deflation increases, and the firm's marginal benefit falls.

The incentive properties of the three policy instruments follow immediately from (3). If $\mathbf{a}$ and p_0 are such that $p_1(\mathbf{a}) < p_0$, then

- Holding s and z fixed, an increase in the matching rate ϕ decreases the abandonment threshold, thus expanding the range of maximum investment;
- Holding ϕ and z fixed, an increase in the baseline subsidy s (which increases the unrestricted portion of the subsidy $s-\alpha z$) increases the abandonment threshold, thus contracting the range of maximum investment;
- Holding ϕ fixed and the unrestricted portion of the subsidy $s-\alpha z$ fixed, an increase in the mandated minimum z increases the abandonment threshold, thus contracting the range of maximum investment.

Suppose, by contrast, $\mathbf{a}$ and p_0 are such that $p_1(\mathbf{a}) \geq p_0$, so there is no range of maximum investment. Then, holding s and z fixed, a sufficiently large increase in the matching rate ϕ could give rise to a

[18] Specifically:

$$\frac{r}{r+\lambda z} = \frac{\frac{E(e^{-r\tilde{T}}|k=z)-E(e^{-r\tilde{T}}|k=1)}{E(e^{-r\tilde{T}}|k=1)}}{\frac{z-1}{1}},$$

where $\widetilde{T}$ denotes the time to discovery; $E(e^{-r\widetilde{T}}|k=1) = \int_0^\infty e^{-rt}\lambda e^{-\lambda t}dt = \frac{\lambda}{\lambda+r}$ is the expected present value of \$1 when the firm invests flat out; and $E(e^{-r\widetilde{T}}|k=z) = \int_0^\infty e^{-rt}\lambda z e^{-\lambda z t}dt = \frac{\lambda z}{\lambda z+r}$ is the expected present value of \$1 when the firm invests at level $z \in [0,1]$. These expressions arise because, conditional on the project being viable, if the investment effort is a constant k, then discovery time is an exponetial random variable with parameter λk.

non-empty range of maximum investment, but increases in either s or z (holding ϕ fixed) would not do so—the firm would continue to invest at the minimum mandated level.

Both the unrestricted component of the subsidy s and the minimum mandate z are a drag on R&D incentives. By investing more heavily and accelerating the expected time to a breakthrough, the firm brings to an end more quickly the flow $s - \alpha z$ of fungible benefits. Increases in unrestricted funding magnify this negative consequence. Furthermore, the minimum mandate z itself becomes a drag on R&D incentives for a different reason. When $k = z$, the firm, in effect, receives a fully-funded option from the government: the government is paying for the R&D investment z, but the firm receives the prize Π if a discovery is made. A firm faces a trade-off between accelerating the time to breakthrough at its own cost and retaining the free option from the government. A larger z increases the option value and thus reduces a firm's incentives to use its own resources for R&D.

Even though an increase in the matching rate expands the range of maximum investment, its impact on cumulative R&D investment is not necessarily "cost effective." We can see this by considering a setting in which the prior p_0 is extremely close to 1. If $p_1(0,0,0) < 1$, an unsubsidized firm would initially invest in R&D, and since the posterior p would initially evolve very slowly (because the firm is virtually certain that the project is viable), a long period of time would have to pass with no breakthrough before the firm abandoned its efforts. A matching subsidy would indeed expand the time the firm persisted with the project, but for a significant period of time the government would be reimbursing some of the firm's R&D expenses even though the firm would have invested in R&D even had it not received that reimbursement over this time frame. A consideration of this sort becomes potentially relevant if the financing of subsidies entails a positive shadow cost of public funds.

3.2 R&D Competition Between Multiple Firms ($N > 1$)

We now turn to the case in which $N > 1$ firms compete to achieve the R&D breakthrough. For this part, we deliberately confine our analysis to symmetric Markov perfect equilibrium (MPE), particularly because its equilibrium refinement implication fits our purpose. The Markov requirement imposed by the equilibrium concept methodologically rules out the possibility that firms tacitly cooperate to overcome the free-rider problem in the presence of R&D spillovers, e.g., using trigger strategies.[19] Therefore,

[19] Trigger strategies depend on state variables which are not intrinsically payoff-relevant. They become payoff-relevant for each player only because the opponents' strategies depend on them. As a refinement of perfect Bayesian equilibria (PBE), an MPE rules out such trigger strategies. The results in Hörner, Klein and Rady (2014) suggest that in our

in the symmetric MPE outcome, we can explore how the free-rider problem interacts with the R&D subsidization. We defer exploring the issue of R&D cooperation to Section 4.4.4, and thus explicitly isolate the R&D cooperation policy from the R&D subsidy policy for the time being. The derivation of the equilibrium follows the approach in Section III of Besanko and Wu (2013).[20] Given a subsidy policy $\mathbf{a}$, there is a unique, symmetric Markov perfect equilibrium (MPE) in the R&D competition game. It is described by an abandonment threshold

$$p_N(\mathbf{a}) = \frac{\alpha\left(1-\phi\right)}{\lambda\left(\left[\Pi - \frac{(s-\alpha z)}{r}\right]\left(\frac{r}{r+\lambda N z}\right) + \left(N-1\right)\left(1-\theta\right)\Pi\left(\frac{\lambda z}{r+\lambda N z}\right)\right)}, \tag{4}$$

a critical spillover level

$$\overline{\theta}(z,s) \triangleq 1 - \frac{\Pi - \left(\frac{s-\alpha z}{r}\right)}{\Pi}\frac{r}{r+\lambda z} \in [0,1],$$

and a slowdown threshold $q_N(\mathbf{a})$ to be defined presently.

If $\theta > \overline{\theta}(z,s)$, the MPE strategy $k_N(p)$ for each firm is

$$k_N\left(p\right) = \begin{cases} 1 & \text{if } p \geq q_N(\mathbf{a}) \\ \frac{r\left[\frac{\lambda\Pi-(1-\phi)\alpha}{\lambda} - B_M(\mathbf{a})(1-p) + \frac{(1-\phi)\alpha(1-p)}{\lambda}\ln\frac{1-p}{p}\right] - s + \alpha\phi z}{(N-1)(\alpha(1-\phi)-\lambda p(1-\theta)\Pi)} & \text{if } p_N(\mathbf{a}) < p < q_N(\mathbf{a}) \\ z & \text{if } p \leq p_N(\mathbf{a}) \end{cases},$$

where $q_N(\mathbf{a})$ and $B_M(\mathbf{a})$ are the values of q_N and B_M that solve the following pair of equations:

$$\frac{r\left[\frac{\lambda\Pi-(1-\phi)\alpha}{\lambda} - B_M\left(1-q_N\right) + \frac{(1-\phi)\alpha(1-q_N)}{\lambda}\ln\left(\frac{1-q_N}{q_N}\right)\right] - s + \alpha\phi z}{\left(N-1\right)\left[\alpha\left(1-\phi\right) - \lambda q_N\left(1-\theta\right)\Pi\right]} = 1, \tag{5}$$

$$\begin{aligned} &\frac{\lambda\Pi-(1-\phi)\alpha}{\lambda} - B_M\left(1-p_N(\mathbf{a})\right) + \frac{(1-\phi)\alpha\left(1-p_N(\mathbf{a})\right)}{\lambda}\ln\left(\frac{1-p_N(\mathbf{a})}{p_N(\mathbf{a})}\right) \\ &= \frac{s-\alpha z}{r} + \frac{\lambda N z\left[\frac{r(1-\theta)\Pi}{N} + r\theta\Pi - s + \alpha z\right]}{r\left(r+\lambda N z\right)}p_N(\mathbf{a}) \end{aligned} \tag{6}$$

setting, non-Markovian equilibria allow the firms to mitigate the free-rider problem to some extent, but not to overcome it completely.

[20] The Online Appendix provides a complete derivation.

If, by contrast, $\theta \leq \overline{\theta}(z,s)$, the MPE strategy for each firm is

$$k_N(p) = \begin{cases} 1 & \text{if } p > p_N(\mathbf{a}) \\ z & \text{if } p \leq p_N(\mathbf{a}) \end{cases}.$$

Panel (a) of Figure 1 illustrates the MPE investment strategy when $\theta > \overline{\theta}(z,s)$. In contrast to monopoly, in which the equilibrium investment policy is "bang-bang," there is a range of beliefs $(p_N(\mathbf{a}), q_N(\mathbf{a}))$ over which $k_N(p) \in (z,1)$. If the posterior belief falls to the slowdown threshold $q_N(\mathbf{a})$, firms start to taper their research efforts by reducing k below 1. If the posterior reaches the abandonment threshold $p_N(\mathbf{a})$, a firm chooses the minimum required R&D effort, $k_N(p) = z$.

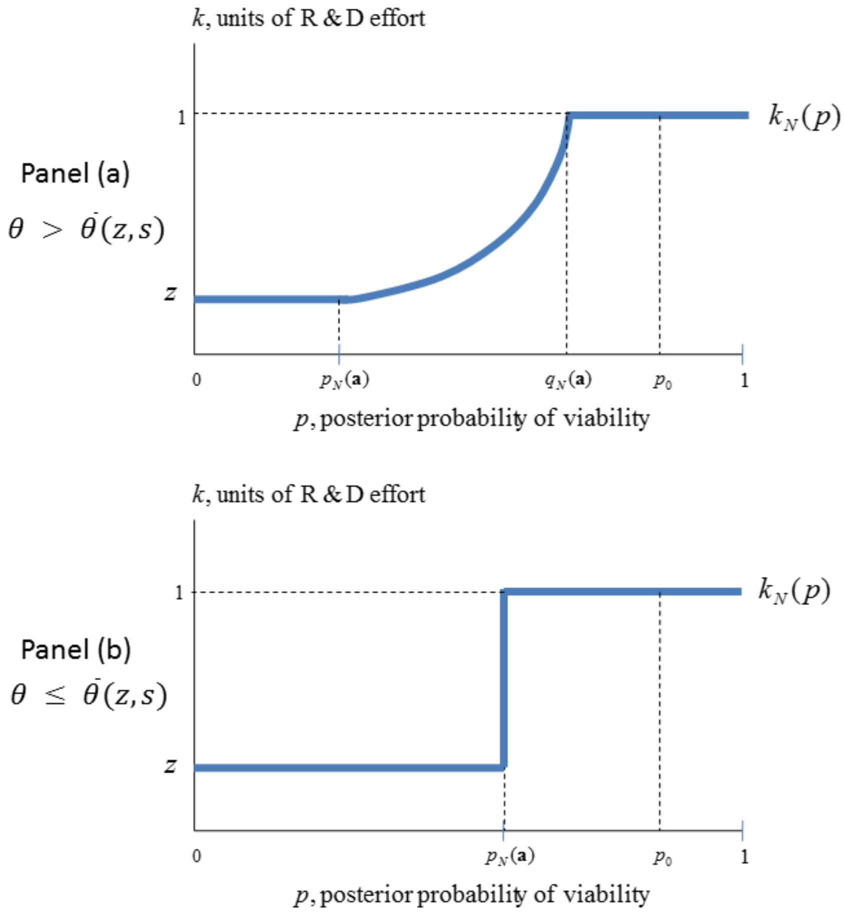

Figure 1: Equilibrium investment strategy with N firms

The equilibrium involves $k_N(p) \in (z, 1)$ because of a free-rider problem: a firm can achieve a positive payoff $\theta\Pi$ from spillovers even if it loses the R&D competition. In particular, when $\theta > \overline{\theta}(z, s)$ and $p \in (p_N(\mathbf{a}), q_N(\mathbf{a}))$, given that all other firms invest "flat out," it will be optimal for a firm to reduce its R&D investment below the maximum level. On the other hand, though, given that all other firms invest at the minimum level, it will be optimal for a firm to invest "flat out" in R&D. The "concession" to the free-rider problem that is made in equilibrium is that for $p \in (p_N(\mathbf{a}), q_N(\mathbf{a}))$, all firms reduce k to a positive number less than 1. The equilibrium value of $k_N(p)$ is such that when a firm's $N - 1$ competitors invest $k_N(p)$, it is indifferent in investing among all $k \in (z, 1)$ and thus chooses $k_N(p)$ in a symmetric equilibrium.

By contrast, when the spillover parameter is less than $\overline{\theta}(z, s)$, the free-rider problem does not arise, and as shown in panel (b) of Figure 1, the equilibrium investment policy is "bang-bang," as in the case of a monopoly. However, the abandonment threshold $p_N(\mathbf{a})$ may not correspond to the monopoly threshold $p_1(\mathbf{a})$. In fact, for any subsidy policy such that $z > 0$, then using (3) and (4), we have[21]

$$\frac{p_1(\mathbf{a})}{p_N(\mathbf{a})} = \frac{\lambda\left(\left[\Pi - \frac{(s-\alpha z)}{r}\right]\frac{r}{r+\lambda N z} + (N-1)(1-\theta)\,\Pi\frac{\lambda z}{r+\lambda N z}\right)}{\lambda\left[\Pi - \left(\frac{s-\alpha z}{r}\right)\right]\frac{r}{r+\lambda z}} > 1, \tag{7}$$

i.e., the range of maximum investment is greater with N firms than with a monopolist. The intuition is this. The fully-funded option that z provides when a firm ceases R&D investment above the minimum mandate becomes less valuable when it competes with $N-1$ other firms, each of whom could potentially get the prize from the breakthrough. As a result, the minimum mandate z is less of a drag on R&D incentives with N firms than with a single firm.

Because $\overline{\theta}$ depends on z and s, whether or not free riding arises in equilibrium depends on the subsidy policy. We note that if there is neither unrestricted funding nor minimum investment mandate — i.e., $z = s = 0$ — then $\overline{\theta}(z, s) = 0$, and the free-rider problem always arises. This tells us that a necessary condition for avoiding the free-rider problem (in the absence of an explicit or tacit R&D consortium) is to establish a mandated minimum level of R&D ($z > 0$) or provide positive baseline funding ($s > 0$), or both. To understand why, recall from the discussion of monopoly that a subsidy with an unrestricted component $s - \alpha z$ creates an implicit loss for the firm when it achieves a breakthrough. By the same token, an unrestricted component creates an implicit loss when *another firm* achieves a breakthrough. Thus, the unrestricted component of the subsidy offsets part of the gain the firm receives when another firm makes the discovery, thereby reducing the temptation to free ride. A subsidy policy with a minimum mandate also creates an implicit loss for the firm when another firm achieves the breakthrough, but for a different reason. Recall that the minimum mandate z creates a fully-funded option for each firm. When another firm achieves a breakthrough, this option goes away, creating an implicit loss that offsets some of the gains from free riding. Thus, in contrast to the monopoly case, in which increases in s and z had unambiguously adverse effects on the provision of R&D, under R&D competition s and/or z may have potentially beneficial incentive effects by mitigating the extent of free riding behavior by firms.

Each of the policy choices affects investment incentives through the entire equilibrium strategy $k_N(p)$. These effects cannot be determined analytically, but the expression for the abandonment threshold in (4) has these implications:

[21] Recall that $\Pi > \frac{(s-\alpha z)}{r}$ given our parameter assumptions, so the condition in (7) is meaningful.

- Holding s and z fixed, an increase in the matching rate ϕ decreases the N-firm equilibrium abandonment threshold, thus expanding the range over which the firms invest in excess of the mandated minimum;

- Holding ϕ and z fixed, an increase in the baseline subsidy s (which thus increases the unrestricted component of the subsidy $s - \alpha z$) increases the N-firm equilibrium abandonment threshold, thus contracting the range over which the firm invests in excess of the mandated minimum.

This discussion hints at an interesting tension involving the baseline subsidy s. On the one hand, it can crowd out private R&D investment, by contracting the range over which the firm invests in excess of the mandated minimum. On the other hand, it can counteract the free-rider problem. As we show below, when the shadow cost of public funds is 0, the government can exploit this tension.

Summarizing the impact of subsidies in the N-firm case, if s and z are sufficiently large, the free-rider problem will not arise in equilibrium. However, changes in ϕ have no impact on free riding. As in the monopoly case, increases in ϕ decrease the abandonment threshold $p_N(\mathbf{a})$ (thus expanding the range of private funding of R&D), while increases in s increase the abandonment threshold. Unlike the monopoly case, changes in z have an ambiguous impact on the abandonment threshold. The impact of z, s, and ϕ on the slowdown threshold $q_N(\mathbf{a})$ and the equilibrium investment policy $k_N(p)$ more generally cannot be determined analytically.

The lesson of this section is that the way in which R&D is subsidized matters. Certain types of subsidies (e.g., earmarked subsidies in the case of $N = 1$) may crowd out private investment, while other types of subsidies (e.g., a pure matching subsidy) may stimulate private investment. This suggests that empirical studies of the impact of R&D subsidies on private R&D investment need to be cognizant of the subsidy mechanism.

4 Optimal Subsidy Policy Under "If" and "When" Uncertainty For an Informationally Constrained Policy Maker

The previous section showed that incentives to invest in R&D depend on how R&D subsidies are structured. We now turn to the question of the best subsidy policy. Recall our assumption that the government neither knows the prior belief p_0 nor has probabilistic beliefs describing its likelihood, which prevents the government from offering a contract that ties the parameters of the subsidy with the posterior belief that determines a firm's R&D strategy.

4.1 First-best R&D Investment Policy

As a benchmark, the first-best R&D investment policy is[22]

$$k^*(p) = \begin{cases} 1 & \text{if } p \in [p^*, 1] \\ 0 & \text{if } p \in [0, p^*] \end{cases}, \tag{8}$$

where

$$p^* = \frac{\alpha}{\lambda\left[CS + \Pi + (N-1)\theta\Pi\right]} < 1. \tag{9}$$

The first-best abandonment threshold p^* is the reciprocal of the social benefit-cost ratio $\frac{\lambda(CS+\Pi+(N-1)\theta\Pi)}{\alpha}$. Note that $p^* < \frac{\alpha}{\lambda\Pi} = p_N(0,0,0)$, so under the first-best policy, investment in the R&D project continues at least as long as it would have in unsubsidized firms, and strictly longer if $p^* < p_0$, i.e., whenever there is any investment in the first-best solution. If $p_0 \in \left(p^*, \frac{\alpha}{\lambda\Pi}\right)$, the first-best policy entails investment for some length of time, while unsubsidized firms opt for no investment at all. Thus, unsubsidized firms underinvest relative to the first-best level.

If the government knew the firm's prior p_0 (and thus the subsequent posterior beliefs p), it could direct the firm to follow the policy in (8) and (9) and achieve the first-best welfare level. Our assumption about the government not knowing p_0 rules this out. Instead, we assume that the government must rely on subsidies to provide incentives to the firm. A subsidy S to the firm entails a transfer of S from taxpayers plus a social cost γS, where $\gamma \geq 0$ is the shadow cost of public funds.

In addition to having a social cost, a subsidy policy could also be distortionary in the sense that it could induce underinvestment or overinvestment in R&D relative to the first-best level. Consider, for example, the case of monopoly R&D and a policy $\mathbf{a}$ such that $p_1(\mathbf{a}) < p^*$ (which would be due to an especially generous matching rate). If $p_0 \in [p_1(\mathbf{a}), p^*]$, the first-best policy calls for no investment at all, while the subsidy policy induces the firm to invest "flat-out" for at least some period of time, and if $z > 0$, the firm would continue to invest once its posterior fell to $p_1(\mathbf{a})$. In this case, we have

[22] Derivation of this policy is a straightforward extension of the derivation of the firm's optimal policy under monopoly R&D presented in the Online Appendix. The extension involves seting $s = z = \phi = 0$, replacing Π with $CS+\Pi+(N-1)\theta\Pi$, and using the transformations $\alpha' = N\alpha$ and $\lambda' = N\lambda$.

overinvestment. On the other hand, for a subsidy policy such that $p^* < p_1(\mathbf{a})$, if $p_0 \in [p^*, p_1(\mathbf{a})]$, the firm would spend none of its own funds investing (though if $z > 0$, it would invest at the minimum mandated level), while the first-best solution would have given rise to non-empty maximum investment range. In this case, then, we would have underinvestment.

4.2 Social Welfare Function

What subsidy policy would the government choose in the face of incomplete information about the prior p_0? To answer this question, we need to derive the expected social welfare induced by the equilibrium investment rules $k_1(p)$ and $k_N(p)$. We describe this derivation for the case of R&D competition; the derivations for the case of monopoly R&D is analogous and thus omitted.

When firms invest $k_N(p)$ units in the R&D project for the time interval $[t, t+dt)$ when their belief about the project's viability is p, the social welfare schedule $W_N(p)$ is given recursively as follows:

$$\begin{aligned} W_N(p) = &-\alpha N k_N(p)dt - \gamma N \left[s + \phi\alpha(k_N(p) - z)\right] dt \\ &+ \lambda N k_N(p) p dt \left[CS + \Pi + (N-1)\theta\Pi\right] + (1 - \lambda N k_N(p) p dt) e^{-rdt} W_N(p + dp). \end{aligned} \tag{10}$$

This can be transformed into a differential equation:

$$\begin{aligned} 0 = &-\alpha N k_N(p) - \gamma N \left[s + \phi\alpha(k_N(p) - z)\right] + \lambda N k_N(p) p \left[CS + (1 + (N-1)\theta)\Pi\right] \\ &- (r + \lambda N k_N(p) p) W_N(p) - \lambda N k_N(p) p (1-p) W_N'(p). \end{aligned} \tag{11}$$

When $k_N(p) = 1$ (i.e., $p > q_N(\mathbf{a})$), the solution to this differential equation is

$$\begin{aligned} W_N(p) = &-\frac{\alpha N}{r} - \frac{N\gamma\left[s + \alpha\phi(1-z)\right]}{r} \\ &+ \left\{\frac{\lambda N\left(r\left[CS + (1 + (N-1)\theta)\Pi\right] + N\alpha + N\gamma\left[s + \alpha\phi(1-z)\right]\right)}{r(r + \lambda N)}\right\} p + B_W^H p \left(\frac{1-p}{p}\right)^{\frac{r+\lambda N}{\lambda N}}, \end{aligned} \tag{12}$$

where B_W^H is a constant. When $k(p) = z$ (i.e., $p < p_N(\mathbf{a})$), the solution to this differential equation is

$$W_N(p) = -\frac{N\alpha z}{r} - \frac{N\gamma s}{r} + \frac{\lambda N z\left(r\left[CS + (1 + (N-1)\theta)\Pi\right] + \alpha N z + N s\gamma\right)}{r(r + \lambda N z)} p.$$

When $k_N(p) \in (z, 1)$, the differential equation (11) does not have a closed-form solution.[23]

The function $W_N(p)$ represents expected social welfare for any arbitrary posterior belief p when firms follow the symmetric equilibrium investment strategy induced by subsidy policy $\mathbf{a}$. It follows that evaluating $W_N(p)$ at $p = p_0$ gives us the government's expected social welfare if it offered policy $\mathbf{a}$ and *it knew the firm's prior belief* p_0. We call this *conditional expected social welfare* (i.e., conditioned on p_0), and we denote it by $W_N(p_0|\mathbf{a})$.

We emphasize that the government cannot actually condition subsidy policy on p_0 (since it does not know what it is), nor can it base policy on an expectation of $W_N(p_0|\mathbf{a})$ with respect to a particular prior distribution over p_0 since we have assumed it does not have a unique prior. That said, it is useful to consider the problem of choosing a subsidy policy to maximize conditional expected welfare:

$$\max_{\mathbf{a} \in A} W_N(p_0|\mathbf{a}). \tag{13}$$

Let $\mathbf{a}^{**}(p_0)$ denote the solution to this problem (and also for the analogous optimization problem under monopoly R&D). If it was the case that the solution to this problem was independent of p_0—i.e., $\mathbf{a}^{**}(p_0) = \mathbf{a}^{**}$—then the government could implement $\mathbf{a}^{**}$ without requiring knowledge of the firm's private information, and it would be assured that it would attain the highest welfare it could possibly achieve whatever the firms' actual type p_0 might be. In this case, the policy in question qualifies as an *ex post* equilibrium. Informally, an *ex post* equilibrium is such that no uninformed player will find it profitable to unilaterally deviate from her strategy even if her incomplete information (about opponents' types) is replaced by complete information. It is obvious to see that here, even if the government knows the true value of p_0, *ceteris paribus* $\mathbf{a}^{**}(p_0)$ would remain an optimal policy. The *ex post* equilibrium status attached to this policy gives it a robust justification in terms of information requirements.

4.3 Optimal R&D Policy Under Incomplete Information: No Shadow Cost of Public Funds ($\gamma = 0$)

We now show that when $\gamma = 0$—a subsidy is a pure transfer between taxpayers and the firm—an *ex post* equilibrium exists in each of two cases considered: monopoly R&D and R&D competition. In each case, the optimal subsidy policy is remarkably simple, and it attains first-best welfare.

[23] In the Online Appendix we derive the social welfare schedule $W_1(p)$ for the special case of $N = 1$. This schedule has a closed-form solution.

Proposition 1 *Under monopoly R&D, if a subsidy is a pure transfer between consumers and firms, i.e.,* $\gamma = 0$, *then the following subsidy policy constitutes an ex post equilibrium:* $\mathbf{a}^{**}(p_0) = \mathbf{a}^{**} = (0, 0, \phi_1^{**})$ *for all* $p_0 \in [0, 1]$, *where*

$$\phi_1^{**} \triangleq \frac{CS}{CS + \Pi} = 1 - \rho.$$

This policy induces the firm to choose an investment policy $k_1(p) = k^*(p)$ *and thus achieves the first-best level of ex ante welfare for any prior belief* p_0.

When subsidies have no shadow cost, there is a simple way to achieve the first-best outcome under monopoly R&D: reimburse the firm a share of its R&D expenditures equal to 1 minus the appropriability ratio, i.e., the consumer surplus share of the social benefit from the breakthrough. Neither a mandated minimum level of R&D nor an unrestricted subsidy is necessary.

When there is no shadow cost of public funds, the first-best outcome can also be attained as an *ex post* equilibrium under R&D competition with multiple firms.

Proposition 2 *Under R&D competition between* $N > 1$ *firms, if a subsidy is a pure transfer between consumers and firms, i.e.,* $\gamma = 0$, *then the following subsidy policy constitutes an ex post equilibrium:* $\mathbf{a}^{**}(p_0) = \mathbf{a}^{**} = (0, s_N^{**}, \phi_N^{**})$ *for all* $p_0 \in [0, 1]$, *where*

$$s_N^{**} \triangleq r\theta\Pi.$$

$$\phi_N^{**} \triangleq \frac{CS + N\theta\Pi}{CS + \Pi + (N-1)\theta\Pi}.$$

This policy induces each firm to choose an investment policy $k_N(p) = k^*(p)$ *and thus achieves the first-best level of ex ante welfare for any prior belief* p_0. *If there are no spillovers* ($\theta = 0$), *the optimal subsidy policy for* N *firms is identical to that for a monopoly.*

With no shadow cost of public funds, the subsidy policy that implements the first-best solution under R&D competition between identical firms is also quite simple. First, the firm receives an unrestricted subsidy s that equals the flow equivalent of the spillover benefits $\theta\Pi$ that it would have received had another firm won the R&D competition. This ensures that the only way that a firm can improve its payoff is by winning the R&D competition, thus eliminating the free-rider problem and making the firms focus on winning the competition. Though a positive value of s eliminates the free-rider problem, it does not fully align the private marginal benefit of R&D with the social marginal benefit of R&D. Through use of a matching rate ϕ equal to $\frac{CS+N\theta\Pi}{CS+\Pi+(N-1)\theta\Pi}$, private and social incentives are aligned.

The comparative statics results that follow from Proposition 2 indicate that under plausible conditions, as the number of competing firms increases, the optimal subsidy mechanism relies more heavily on the matching rate to align private and social incentives.

Corollary 1 *Suppose that the payoff* $\Pi(N)$ *from the breakthrough is non-increasing in the number of firms, i.e.,* $\frac{d\Pi}{dN} \leq 0$*, the consumer surplus* $CS(N)$ *is non-decreasing in* N*, i.e.,* $\frac{dCS}{dN} \geq 0$*, and spillovers are strictly positive but less than complete, i.e.,* $\theta \in (0,1)$*. Under R&D competition between* $N > 1$ *firms and if a subsidy is a pure transfer between consumers and firms: (a) the unrestricted component* s_N^{**} *of the optimal subsidy is non-increasing in the number of firms* N*, while the optimal matching rate* ϕ_N^{**} *is strictly increasing in* N*; (b) if* $\Pi(N)$ *and* $CS(N)$ *are independent of* N*, then as the number of competing firms increases without bound,* s_N^{**} *remains unchanged, while* $\phi_N^{**} \to 1$*; (c) if* $\Pi(N)$ *strictly decreases in* N *and, in addition,* $\lim_{N\to\infty} \Pi(N) = 0$*,* $\lim_{N\to\infty} N\Pi(N) < \infty$ *and* $CS(1) > 0$*, then as the number of competing firms increases without bound,* $s_N^{**} \to 0$*, while* $\phi_N^{**} \to 1$*.*

The implication that the matching rate increases as the number of firms increases illustrates that in our setting more competition *per se* does not serve to align private and social incentives.

4.4 Optimal R&D Policy Under Incomplete Information: Positive Shadow Cost of Public Funds ($\gamma > 0$)

4.4.1 Non-existence of an *Ex-Post* Equilibrium

The results in the preceding section do not extend to the case in which the shadow cost of public funds γ is positive. The first-best policy cannot be implemented through subsidies, and indeed, no *ex post* equilibrium exits for $\gamma > 0$. We first establish this for the case of monopoly R&D and then for non-cooperative R&D.

Proposition 3 *Under monopoly R&D, if* $p_1(0,0,0) = \frac{\alpha}{\lambda\Pi} < 1$*, then for* $\gamma > 0$ *there exists no ex post equilibrium, i.e., the solution to* $\max_{\mathbf{a}\in A} W_1(p_0|\mathbf{a})$ *depends on* p_0*.*[24]

[24] Numerical examples show that the restriction $p_1(0,0,0) < 1$ (which is useful for a simple analytical proof) is unnecessary for the non-existence result.

The intuition for this result is as follows. When the shadow cost of public funds is positive, the government faces a trade-off between using a subsidy to induce more R&D and incurring higher social costs due to the subsidy, and this trade off delicately depends on the value of p_0. To see why, suppose that the government *knew* that the firm was completely certain of the project's viability, i.e., $p_0 = 1$. In this case, the subsidy policy that solves the optimization in (13) is no subsidy at all. This is because if $p_1(0,0,0) < 1$ the government would be certain that an unsubsidized firm would invest in R&D, no matter how much time passed without a breakthrough. Recall that when $p_0 = 1$, there is no "if" uncertainty, and thus $p(t) = 1$ for all t. Thus, a socially costly subsidy would have no impact on the firm's investment incentives. The problem, however, is that the no-subsidy policy $\mathbf{a} = (0,0,0)$, while maximizing welfare if the government was certain that $p_0 = 1$, could be a poor policy if the government believed that the firm's prior was something else.

A non-existence result also arises under R&D competition.

Proposition 4 *Under R&D competition with $N > 1$ firms, if $p_N(0,0,0) < q_N(0,0,0) < 1$, then for $\gamma > 0$ there exists no ex post equilibrium.*[25]

The intuition for this result is essentially the same as that for Proposition 3: the trade-off between the higher R&D investment induced by a subsidy and the higher social cost resulting from the subsidy depends on the value of p_0 that is unknown to the government.

The non-existence problem limits the general applicability of the solution concept of *ex post* equilibrium. For this reason, we now consider an alternative decision criterion for the government when $\gamma > 0$: max-min.

4.4.2 Max-Min Subsidy Policy Under Ambiguity With a Positive Shadow Cost of Public Funds

In this section, we consider the implications of using max-min as the government's decision criterion. To fix ideas, we begin with the case of monopoly R&D, using it to explain why max-min is an appealing decision criterion. We then apply the max-min criterion to R&D competition with N firms. But as prelude to our discussion of max-min, we show that in the case of monopoly R&D there exists a simple

[25] The restrictive condition: $p_N(0,0,0) < q_N(0,0,0) < 1$ is imposed for the convenience of an analytical proof. Numerical examples show it is not a necessary condition for the result to hold.

policy with the appealing feature that it neither generates underinvestment nor overinvestment. Because this particular policy plays an important role in the characterization of max-min subsidy policies, we begin there.

Solving the Problem of Underinvestment without Inducing Overinvestment We saw earlier that an unsubsidized firm has a tendency to underinvest (relative to the first best), but that by subsidizing the firm, the government could conceivably induce overinvestment. With the government lacking the ability to make fine-tuned trade-offs between policy instruments because it has neither complete nor probabilistic information about p_0, a plausible criterion for a "good" R&D subsidy policy would be one that solves the problem of underinvestment, while not inducing overinvestment. In this section, we show there is such a policy. Indeed, that policy is a natural extension of the policy characterized in Proposition 1, and it smoothly approaches that policy in the limit as $\gamma \to 0$.

Consider a matching rate

$$\phi_{\gamma 1}^{**} \triangleq \frac{\frac{CS}{1+\gamma}}{\frac{CS}{1+\gamma} + \Pi}. \tag{14}$$

Clearly, $\phi_{\gamma 1}^{**} = \phi_1^{**} = 1 - \rho$ for $\gamma = 0$. Further, define the abandonment threshold under the subsidy policy $\left(0, 0, \phi_{\gamma 1}^{**}\right)$ as

$$p_1\left(0, 0, \phi_{\gamma 1}^{**}\right) = p_{\gamma 1}^{**} \triangleq \frac{\alpha}{\lambda\left(\Pi + \frac{CS}{1+\gamma}\right)}. \tag{15}$$

We now note two key properties of the policy $\left(0, 0, \phi_{\gamma 1}^{**}\right)$:

$$p\lambda\left(\Pi + CS\right) - \alpha - \gamma\phi_{\gamma 1}^{**}\alpha \geq 0 \text{ iff } p \geq p_{\gamma 1}^{**}, \tag{16}$$

and

$$p\lambda\Pi - \left(1 - \phi_{\gamma 1}^{**}\right)\alpha \geq 0 \text{ iff } p \geq p_{\gamma 1}^{**}. \tag{17}$$

Note that the left-hand sides of the inequalities in (16) and (17) are the net flow payoffs for the government and the firm, respectively. Therefore, under subsidy policy $\left(0, 0, \phi_{\gamma 1}^{**}\right)$ both the firm and the government have zero net flow payoffs for all $k \in [0, 1]$ for $p = p_{\gamma 1}^{**}$, and their net flow payoffs strictly increase in p. Note that the two zero net flow payoff conditions imply that the "marginal cost equal to marginal benefit condition" is satisfied at the same $p_{\gamma 1}^{**}$ for both the government and the firm. In addition, both the firm and the government have zero continuation values for $p = p^{**}_{\gamma 1}$. Thus, from the government's perspective, given that it has committed to pay a matching rate $\phi^{**}_{\gamma 1}$, $p^{**}_{\gamma 1}$ is the optimal abandonment threshold. And from the firm's perspective, given that it receives a pure matching subsidy $\phi^{**}_{\gamma 1}$, $p^{**}_{\gamma 1}$ is the privately optimal abandonment threshold. Thus, if the government offers the policy $\left(0, 0, \phi^{**}_{\gamma 1}\right)$, it can be assured that whatever the firm's private information it will invest in the way most preferred by the planner.

We note that the policy $\left(0, 0, \phi^{**}_{\gamma 1}\right)$ has the intuitive feature that the matching rate is smaller than ϕ^{**}_{1}, and this rate decreases as the shadow cost of public funds increases. As the shadow cost increases without bound, the matching rate goes to zero, i.e., the policy involves no subsidization whatsoever.

Max-Min Policies: Monopoly R&D ($N = 1$) We now turn to the max-min criterion. This criterion has been featured prominently in the literature on decision under ambiguity (for example, see Gilboa and Schmeidler, 1989). Ambiguity can, in general, be modeled as a set of multiple probabilistic beliefs, allowing the flexibility to model belief-free incomplete information games, to which the concept of max-min solution can be applied. In a policy making context, max-min would correspond to a setting in which, although the government does not know *what to believe* about the firm's private information, it wishes to follow a policy whose "worst case scenario" outcome is not inferior to any alternative policy. Max-min is a potentially appealing decision criterion when policy makers (as often seems plausible) are especially attuned to the need to avoid potentially severe policy mistakes. We begin with a general definition of max-min subsidy policies, allowing for the possibility of mixed strategies.

Definition 1 *Let Σ_A denote the Borel sigma algebra over A, and $\Delta\left(\Sigma_A\right)$ the set of all probability measures (i.e., mixed strategies) over Σ_A. Let Σ be the Borel sigma algebra over $[0,1]$ and $\Delta\left(\Sigma\right)$ the set of all probability measures (i.e., probabilistic beliefs) over Σ. A mixed strategy $\sigma \in \Delta\left(\Sigma_A\right)$ is a max-min solution if*

$$\sigma = \arg \max_{\sigma' \in \Delta(\Sigma_A)} \min_{\mu \in \Delta(\Sigma)} \int_0^1 \int_{\mathbf{a} \in A} W_1\left(p|\mathbf{a}\right) d\sigma'\left(\mathbf{a}\right) d\mu\left(p\right).$$

An alternative decision criterion for belief-free incomplete information games is rationalizability (Pearce, 1984; Battigalli and Sciniscalchi, 2003). In the Online Appendix we explore rationalizability as an alternative to max-min. We show that while all max-min solutions in our context are rationalizable, the converse does not hold. That is, not all rationalizable policies are max-min solutions. This implies that the max-min concept has more bite than rationalizability in our context, providing some justification for our focus on max-min as a decision criterion.

We now characterize the set of all pure strategy max-min subsidy policies under monopoly R&D when $\gamma > 0$.

Proposition 5 *Suppose $\gamma > 0$. Under monopoly R&D, the max-min value of all possible (mixed strategy) subsidy policies is zero, i.e.,*

$$\max_{\sigma' \in \Delta(\Sigma_A)} \min_{\mu \in \Delta(\Sigma)} \int_0^1 \int_{\mathbf{a} \in A} W_1 \left(p|\mathbf{a}\right) d\sigma' \left(\mathbf{a}\right) d\mu \left(p\right) = 0.$$

Moreover, the set of all pure-strategy max-min solutions is the set of pure matching subsidy policies given by

$$A_{1\max\min} \triangleq \left\{(0, 0, \phi) \,|\phi \in \left[0, \phi^{**}_{\gamma 1}\right]\right\}.$$

The set of all mixed-strategy max-min solutions is

$$\left\{\sigma \in \Delta\left(\Sigma_A\right) |\sigma \textit{ is a mixed strategy over support } A_{1\max\min}\right\}.$$

Proposition 5 implies that pure strategy max-min subsidy policies do not entail a minimum mandate or an unrestricted subsidy component, and they have matching rates ranging from 0 (no subsidy whatsoever) up to $\phi^{**}_{\gamma 1}$, the matching subsidy that induces neither over- nor underinvestment.

To build intuition for this result, consider a policy $\mathbf{a}_z \triangleq (z, \alpha z, 0)$ that involves a minimum mandate. Of course, the government lacks information about the firm's beliefs about project viability and must consider the possibility that p_0 could be low. It is under these circumstances that policy $\mathbf{a}_z$ is unappealing. Using the expression for $W_1\left(p_0|\mathbf{a}\right)$ in equation (24) in the Appendix, we can show that $W_1\left(p_0|\mathbf{a}_z\right) < 0$ for $p_0 \in \left[0, \underline{p}\left(\mathbf{a}_z\right)\right]$, where $\underline{p}\left(\mathbf{a}_z\right)$ is defined such that $W_1\left(\underline{p}\left(\mathbf{a}_z\right)|\mathbf{a}_z\right) \equiv 0$. It follows immediately that $\min_{\mu \in \Delta(\Sigma)} \int_0^1 W_1\left(p|\mathbf{a}_z\right) d\mu\left(p\right) < 0$, and hence $\mathbf{a}_z$ is not a max-min solution. Policy $\mathbf{a}_z$ performs poorly when p_0 is low because, as we have seen, when a firm is sufficiently pessimistic about the possibility of a breakthrough it will only invest at the minimum required level z. In this case, a policy with a minimum mandated level of R&D forces the firm to expend resources on a project that is unlikely to result in a breakthrough, and moreover under policy $\mathbf{a}_z$, it provides the firm with a series of subsidy payments that entail a social cost $\gamma\alpha z$.

Consistent with the literature on public policy mechanisms in the face of ambiguity discussed in the introduction to the paper, one would expect that the max-min policies would reflect a posture of precaution. In our context, a matching subsidy errs on the side of caution because the government pays only for the R&D investment the firm actually makes. Because this investment is correlated with the firm's beliefs about project viability, the government minimizes its risk of making excessive expenditures on socially costly subsidies through reimbursements based on actual investment. Though our model is specific, we believe that this insight about subsidy policy is robust: when funds are socially costly, exercising precaution when subsidizing socially valuable activities should involve tying subsidy payments to observable metrics that tend to lower the payments in "worst case scenarios." Although theoretically this insight seems to follow trivially from the assumption that the government views arbitrarily pessimistic priors (and even $p_0 = 0$) as possible, it is just one manifestation of the deeper insight about context-specific precaution induced by a preference for robustness. Another manifestation is the insight that a precaution against socially inefficient overinvestment leads to a cap on the reimbursement rate of the pure matching subsidy. This result is not driven by the possibility of extremely pessimistic priors; to the contrary, it is based on the possibility of moderately pessimistic priors (e.g., $p_0 = p^{**}_{\gamma 1}$), since these are the circumstances under which a matching subsidy with a high reimbursement rate is prone to inducing overinvestment.

4.4.3 Max-Min Policies: R&D Competition Between Multiple Firms ($N > 1$)

The logic underlying Proposition 5 can be also be applied to the case of $N > 1$. Analogous to the matching rate and the abandonment threshold in (14) and (15), we have

$$\phi^{**}_{\gamma N} \triangleq \frac{\frac{CS+(N-1)\theta\Pi}{1+\gamma}}{\frac{CS+(N-1)\theta\Pi}{1+\gamma} + \Pi} \tag{18}$$

$$p_N\left(0, 0, \phi^{**}_{\gamma N}\right) = p^{**}_{\gamma N} \triangleq \frac{\alpha}{\lambda\left[\Pi + \frac{CS+(N-1)\theta\Pi}{1+\gamma}\right]}. \tag{19}$$

Under the max-min criterion, $s = z = 0$ and the set of max-min subsidy policies consists entirely of pure matching subsidies.

Proposition 6 [26] *Suppose $\gamma > 0$. Under R&D competition between $N > 1$ firms, the set of all pure-*

[26] Because the proof of Proposition 6 follows very similar logic as the proof of Proposition 5, it is omitted.

strategy max-min solutions is the set of pure matching subsidy policies given by

$$A_{N\max\min} \triangleq \left\{(0,0,\phi)\,|\phi \in \left[0,\phi_{\gamma N}^{**}\right]\right\}.$$

Under the max-min decision criterion, if the optimal subsidy policy has a further emphasis on overcoming the possibility of underinvestment, the best policy would be $\left(0,0,\phi_{\gamma N}^{**}\right)$, whose matching rate is less than the optimal rate ϕ_N^{**} in the absence of a shadow cost of public funds. This policy induces investment to terminate at $p = p_{\gamma N}^{**}$ which is the abandonment threshold most preferred by the government given that it has committed to the matching rate $\phi_{\gamma N}^{**}$.[27] However, since $(z,s) = (0,0)$, the threshold $\bar{\theta}(z,s) = 0 \leq \theta$, (in the absence of explicit or tacit R&D cooperation) this policy cannot overcome the inter-firm free-rider problem, and it therefore cannot avoid underinvestment in terms of investment intensity; i.e., investment is not "flat-out" for $p \in \left(p_{\gamma N}^{**}, q_N\left(0,0,\phi_{\gamma N}^{**}\right)\right)$.

Comparative statics analysis of the highest max-min matching rate $\phi_{\gamma N}^{**}$ yields the following result.

Corollary 2 *Suppose that the payoff $\Pi(N)$ from the breakthrough is non-increasing in the number of firms, i.e., $\frac{d\Pi}{dN} \leq 0$, the consumer surplus $CS(N)$ is non-decreasing in N, i.e., $\frac{dCS}{dN} \geq 0$, and spillovers are strictly positive but less than complete, i.e., $\theta \in (0,1)$. Under R&D competition between $N > 1$ firms and $\gamma > 0$: (a) the highest max-min matching rate $\phi_{\gamma N}^{**}$ is strictly increasing in N; (b) if $\Pi(N)$ and $CS(N)$ are independent of N, then as the number of competing firms increases without bound, $\phi_{\gamma N}^{**} \to 1$; (c) if $\Pi(N)$ strictly decreases in N and, in addition, $\lim_{N\to\infty}\Pi(N) = 0$ and $\lim_{N\to\infty} N\Pi(N) < \infty$ and $CS(1) > 0$, then as the number of competing firms increases without bound, $\phi_{\gamma N}^{**} \to 1$.*[28]

The highest max-min matching rate $\phi_{\gamma N}^{**}$ thus varies with the number of competing firms in the same way as the optimal matching rate ϕ_N^{**} when there is no shadow cost of public funds. In fact, in light of Corollaries 1 and 2, when $\lim_{N\to\infty}\Pi(N) = 0$, the optimal subsidy policy with no shadow cost and the highest max-min subsidy policy with a positive shadow cost converge toward each other as $N \to \infty$.

4.4.4 The Benefits of R&D Cooperation

Because the max-min subsidy policy under R&D competition cannot overcome the inter-firm free-rider problem when there is a positive shadow price of public funds, allowing firms to cooperate in making

[27] This claim is verified by equation (26) in the proof of Proposition 4.

[28] The proof of this corollary follows steps of calculus and algebra analogous to the proof of Corollary 1 and is thus omitted.

R&D decisions may stimulate R&D investment and economize on socially costly subsidies. To explore this, we consider an N-firm research consortium that makes R&D decisions cooperatively.

When faced with subsidy policy $\mathbf{a}$, each firm in the consortium has an optimal R&D policy $k_C(p)$ that is identical to the optimal policy $k_1(p)$ in the monopoly case but with a single firm's prize Π replaced by the consortium's collective prize $\Pi + (N-1)\,\theta\Pi$.[29] That is,

$$k_C(p) = \begin{cases} 1 & \text{if } p > p_N^C(\mathbf{a}) \\ z & \text{if } p \leq p_N^C(\mathbf{a}) \end{cases},$$

where the abandonment threshold $p_C(\mathbf{a})$ is given by

$$p_C(\mathbf{a}) = \frac{\alpha(1-\phi)}{\lambda\left[\Pi + (N-1)\,\theta\Pi - \frac{s-\alpha z}{r}\right]\left(\frac{r}{r+\lambda z}\right)}. \tag{20}$$

It is straightforward to establish (again mimicking the analysis for monopoly R&D) that in the absence of a shadow cost of public funds, first-best welfare can be attained by use of a pure matching subsidy, with a matching rate

$$\phi_C^{**} \triangleq \frac{CS}{CS + \Pi + (N-1)\theta\Pi} = 1 - \rho.$$

Recalling Proposition 2, we observe that when $\theta > 0$, the matching rate ϕ_C^{**} needed to attain the first-best outcome under R&D cooperation is less than the matching rate ϕ_N^{**} needed to attain the first-best outcome under non-cooperative research. Further, attaining the first-best outcome under R&D competition entails an unrestricted subsidy in addition to a matching subsidy. Thus, when there is no shadow cost of public funds, attaining the first-best outcome with a research consortium involves a smaller overall subsidy than attaining the first-best outcome with N non-cooperative firms. This suggests that a research consortium may have an advantage over non-cooperative research when the shadow cost of public funds is positive.

To verify this intuition, we establish that the set of max-min subsidy policies under a research consortium involve pure matching subsidies with a matching rate analogous to that of monopoly.

[29] Further discussion of the derivation of this policy is presented in the Online Appendix.

Proposition 7 [30] *Suppose $\gamma > 0$. Under R&D cooperation through a research consortium with N firms, the set of all pure-strategy max-min solutions is the set of pure matching subsidy policies given by*

$$A_{C\max\min} \triangleq \left\{(0,0,\phi)\,|\phi \in \left[0,\phi^{**}_{\gamma C}\right]\right\},$$

where

$$\phi^{**}_{\gamma C} = \frac{\frac{CS}{1+\gamma}}{\frac{CS}{1+\gamma} + \Pi + (N-1)\theta\Pi}. \tag{21}$$

Under the max-min decision criterion and R&D organized in a research consortium, if the optimal subsidy policy has a further emphasis on overcoming possible underinvestment, then the best policy would be $\left(0,0,\phi^{**}_{\gamma C}\right)$. This policy can overcome the possibility of underinvestment, both in terms of cumulative investment and investment intensity at each point of time. The lowest abandonment threshold is

$$p^{**}_{\gamma C} = \frac{\alpha\left(1-\phi^{**}_{\gamma C}\right)}{\lambda\left[\Pi + (N-1)\theta\Pi\right]} = \frac{\alpha}{\lambda\left[\frac{CS}{1+\gamma} + \Pi + (N-1)\theta\Pi\right]}. \tag{22}$$

Comparing (21) to (18) and (22) to (19) we have

Corollary 3 *Suppose $\gamma > 0$. The highest max-min matching rate under a research consortium is less than the highest max-min matching rate under R&D competition, i.e., $\phi^{**}_{\gamma C} < \phi^{**}_{\gamma N}$, and the corresponding abandonment threshold under a research consortium is less than that under R&D competition, i.e., $p^{**}_{\gamma C} < p^{**}_{\gamma N}$.*

By enabling the formation of a research consortium, a government using the max-min decision criterion can economize on the shadow cost $\gamma\phi\alpha$ of subsidization, while inducing more investment than would occur under R&D competition. Ultimately, this is because the research consortium eliminates the inter-firm free-rider problem that arises under R&D competition. This can be seen most directly by noting that for a given matching rate, a pure matching subsidy under a research consortium results in a smaller abandonment threshold than under R&D competition, i.e.,

$$p_C\left(0,0,\phi\right) = \frac{\alpha(1-\phi)}{\lambda\left[\Pi + (N-1)\,\theta\Pi\right]} < p_N\left(0,0,\phi\right) = \frac{\alpha(1-\phi)}{\lambda\Pi}.$$

[30] The proof of Proposition 7 follows very similar logic as the proof of Proposition 5 and is thus omitted.

5 Conclusions

In this paper, we study the optimal subsidy policy for research programs when the firm is privately informed about project viability and the government is unable to form a unique prior belief about the firm's private information. We first showed that different subsidy tools affect the firm's R&D incentives in different ways. In the case of monopoly R&D, a matching subsidy can stimulate R&D activity, while earmarked and unrestricted subsidies can suppress R&D activity. In the case of R&D competition, the incentive effects of subsidies are somewhat more complex. As in the case of monopoly R&D, an increase (*ceteris paribus*) in the matching rate expands the range over which firms invest in excess of the mandated minimum, while increases in the baseline unrestricted subsidy have the opposite effect. However, R&D competition involves the possibility of free riding, and increases in both the mandated minimum and the unrestricted component of the subsidy can mitigate the extent of free riding by the firms.

We then studied the government's choice of an optimal subsidy policy. If there is no shadow cost of public funding, we show that the government can attain the first-best welfare outcome as an *ex post* (belief-free) equilibrium through the use of very simple subsidy policies. In the case of monopoly, the optimal subsidy scheme is a pure matching policy with the matching rate equal to the ratio of the portion of social welfare not appropriable by the firm to total social welfare. In the case of competition, the optimal policy is a combination of two instruments: a matching subsidy and an unrestricted subsidy. The unrestricted subsidy eliminates firms' incentives to free ride, and the matching subsidy solves the underinvestment due to the appropriability problem. Together, they ensure firms to follow the first-best investment path.

An *ex post* equilibrium does not exist when the shadow cost of public funding is positive. This necessitates consideration of what the government's objective should be in selecting an optimal policy. We consider the max-min criterion. We show that for both $N = 1$ and $N > 1$, the set of max-min policies consists entirely of pure matching subsidies. When $N = 1$, the policy with the highest matching rate solves the problem of underinvestment without inducing overinvestment. This is not the case when $N > 1$, since a pure matching subsidy is unable to correct for underinvestment that arises due to the free-rider problem. However, the highest max-min matching rate for a research consortium is less than the highest max-min matching rate under R&D competition, so allowing firms to determine R&D levels cooperatively can economize on the total shadow cost of the subsidy policy.

There are a number of interesting issues not addressed in this paper that warrant further attention. First, the belief updating structure in our paper is rather simple. As more time passes without a discovery, the updated likelihood that the project is viable falls. The simplicity of this updating rule allows us to derive a closed-form solution to our problem. However, this rule does not allow upward revision of the viability probability. To model this, we need to allow for the possibility that firms acquire new information as the research program progresses. This would be a useful extension of our model.

Another useful extension would be to allow for the possibility of subsidy policies that provide firms with (time-invariant) lump-sum payments contingent on reaching periodic cumulative effort targets. Such subsidies are not uncommon in the real world since, in some settings, a firm may be required to present a prototype or detailed proposal to receive funding. Subsidies contingent on achieving cumulative effort targets could induce firms to exert maximum effort in between payments and thus could potentially economize on socially costly subsidy payments. However, a potential drawback of such schemes in our context of severe informational constraints is that it seems likely that the conditional welfare-maximizing effort targets would depend on the firms' prior belief p_0, with weaker targets and/or larger payments being required for firms that are extremely pessimistic about project viability. It is therefore not clear whether in a max-min subsidy policy contingent payments would be preferred to a pure matching subsidy since the worst-case scenario might entail large contingent subsidies or distortions in effort targets.

Finally, as noted above, our paper restricts attention to time-invariant subsidy policies. In particular, this rules out the use of funding deadlines and terminations as a way of motivating R&D investment. Bonatti and Hörner (2011) study the use of deadlines in addressing the free-riding problem in an R&D collaboration. They show that a finite deadline T can be chosen to induce the agents to contribute maximum effort throughout the process. The intuition is that because agents cannot continue the project on their own once the deadline is reached, as the deadline approaches agents will increase their effort level to race against time. They show that there is an optimal time $T' < T$ after which the agents will exert maximum effort levels. Under certain conditions, T can be chosen so that $T' \leq 0$, which implies that the agents will exert full effort throughout the collaboration process.

Our paper differs from Bonatti and Hörner in that private incentives in our model are not only affected by the free-rider problem, but also by the appropriability problem. While Bonatti and Hörner show that a deadline can neutralize the free-rider problem, it is less clear that a deadline would fully neutralize the appropriability problem. Given the potential consumer benefit from the project that cannot be appropriated by firms, a social planner may not want to set a deadline such that firms are forbidden to conduct research after the deadline expires. Further, in the context of our model, the government could not completely forbid firms to conduct self-funded research, and so a strict deadline on research activity would not be feasible. Still, deadlines on subsidies may be useful in our model. Removing governmental support after a certain point in time would, as in Bonatti and Hörner, generate additional incentives. However, we conjecture that in contrast to the Bonatti and Hörner model, a subsidy deadline by itself would not achieve the first-best outcome when the appropriability problem is present. More generally, though, allowing for the possibility of a time-varying subsidy mechanism may move the social planner closer to the first-best solution in those cases in which time-invariant policies cannot attain the first-best solution (i.e., when there is a shadow-cost of public funds). In such cases, it would be useful to explore the interaction of deadlines and other subsidy instruments and in particular, whether a deadline is a complement to the matching rate, or a substitute for it, in generating incentives.

6 Appendix

Proof of Proposition 1:

To prove the proposition, we will show that the first-best level of welfare can be attained by setting $s = z = 0$ and choosing $\phi = 1 - \rho$. When $N = 1$, the first-best R&D policy solves

$$W^*(p) = \max_{k\in[0,1]} \left[-\alpha k dt + \lambda k p dt \,(CS + \Pi) + (1 - \lambda k p dt)e^{-rdt} W^*(p + dp)\right], \tag{23}$$

and it is given by

$$k^*(p) = \begin{cases} 1 & \text{if } p \in [p^*, 1] \\ 0 & \text{if } p \in [0, p^*] \end{cases},$$

where

$$p^* = \frac{\alpha}{\lambda(\Pi + CS)}.$$

Now, if $s = z = 0$,

$$p_1(\mathbf{a}) = p_1(0, 0, \phi) = \frac{\alpha(1 - \phi)}{\lambda \Pi}.$$

A matching rate given by $\phi = 1 - \frac{\Pi}{\Pi + CS} = 1 - \rho$, along with $s = z = 0$, ensures that $p_1(\mathbf{a}) = p^*$, and thus implements the maximum level of expected welfare for any prior belief p_0. Therefore, $z = 0, s = 0$, $\phi = 1 - \rho$ is the optimal policy.■

Proof of Proposition 2:

We employ the same logic as in the proof of Proposition 1. Recall that under the equilibrium policy there is no free riding if and only if

$$\theta \leq \overline{\theta}(z, s) = 1 - \frac{\Pi - \left(\frac{s - \alpha z}{r}\right)}{\Pi} \frac{r}{r + \lambda z},.$$

When $z = 0$, $s = r\theta\Pi$, then $\overline{\theta}(z, s) = \theta$, which is just enough to eliminate the free-rider problem. The equilibrium investment policy in this case is

$$k_N(p) = \begin{cases} 1 & \text{if } p > p_N(0, r\theta\Pi, \phi) \\ 0 & \text{if } p \leq p_N(0, r\theta\Pi, \phi), \end{cases},$$

where

$$p_N(0, r\theta\Pi, \phi) = \frac{\alpha(1-\phi)}{\lambda(1-\theta)\Pi}.$$

By setting $\phi = \frac{CS+N\theta\Pi}{CS+\Pi+(N-1)\theta\Pi}$, we can make $p_N(0, r\theta\Pi, \phi) = p^* = \frac{\alpha}{\lambda[CS+\Pi+(N-1)\theta\Pi]}$. Thus, the first-best investment policy can be induced by setting $z = 0$ and using a combination of unrestricted funding $s = r\theta\Pi$ and a matching rate $\phi = \frac{CS+N\theta\Pi}{CS+\Pi+(N-1)\theta\Pi}$.■

Proof of Corollary 1:

(a) It is clear that if $\Pi(N)$ is non-increasing in N and $CS(N)$ is non-decreasing in N, the unrestricted component of the subsidy $r\theta\Pi(N)$ is non-increasing in N. For the matching rate, we have

$$\frac{d\phi_N^{**}}{dN} = \frac{-(1-\theta)CS\frac{d\Pi}{dN} + \frac{dCS}{dN}(1-\theta)\Pi + \theta(1-\theta)\Pi^2}{[CS+\Pi+(N-1)\theta\Pi]^2} > 0,$$

since $\frac{d\Pi}{dN} \leq 0$, $\frac{dCS}{dN} \geq 0$ and $\theta \in (0,1)$.

(b) The result follows immediately from the expression for the matching rate ϕ_N^{**}.

(c) When $\Pi(N)$ strictly decreases in N, $\lim_{N\to\infty} \phi_N^{**} = 1$, since $\lim_{N\to\infty} \Pi(N) = 0$, $\lim_{N\to\infty} N\Pi(N) < \infty$ and $\lim_{N\to\infty} CS(N) \geq CS(1) > 0$ by assumption.■

Proof of Proposition 3:

Suppose the contrary, i.e., there exists a policy $\mathbf{a} \in A$ such that $\mathbf{a}$ does not depend on p_0, but maximizes $W_1(p_0|\mathbf{a})$ for all $p_0 \in [0,1]$. From the derivation of the social welfare schedule $W_1(p)$ in the Online Appendix, we have

$$W_1(p_0|\mathbf{a}) = \begin{cases} \Psi(B_W(p_1(\mathbf{a}), \mathbf{a}), \mathbf{a}) & p_0 \in [p_1(\mathbf{a}), 1] \\ -\left(\frac{\alpha z+\gamma s}{r}\right) + \frac{\lambda z}{r+\lambda z}\left[CS+\Pi+\left(\frac{\alpha z+\gamma s}{r}\right)\right]p_0 & p_0 \in [0, p_1(\mathbf{a})] \end{cases}, \tag{24}$$

where

$$\Psi(B_W, \mathbf{a}) \triangleq -\left(\frac{\alpha+\gamma s+\gamma\phi\alpha(1-z)}{r}\right) + \frac{\lambda}{r+\lambda}\left[CS+\Pi+\left(\frac{\alpha+\gamma s+\gamma\phi\alpha(1-z)}{r}\right)\right]p_0 + B_W p_0\left(\frac{1-p_0}{p_0}\right)^{\frac{r+\lambda}{\lambda}},$$

$$B_W(p_1, \mathbf{a}) \triangleq \frac{\frac{\alpha(1-z)}{r} + \frac{\gamma\phi\alpha(1-z)}{r} - \frac{\lambda}{r+\lambda}\left[CS+\Pi+\left(\frac{\alpha+\gamma s+\gamma\phi\alpha(1-z)}{r}\right)\right]p_1 + \frac{\lambda z}{r+\lambda z}\left[CS+\Pi+\left(\frac{\alpha z+\gamma s}{r}\right)\right]p_1}{p_1\left(\frac{1-p_1}{p_1}\right)^{\frac{r+\lambda}{\lambda}}},$$

and

$$p_1 = p_1(\mathbf{a}).$$

Now, when $p_0 = 1$, by (24) we have

$$W_1(1|\mathbf{a}) = -\left(\frac{\alpha + \gamma s + \gamma\phi\alpha(1-z)}{r}\right) + \frac{\lambda}{r+\lambda}\left[CS + \Pi + \left(\frac{\alpha + \gamma s + \gamma\phi\alpha(1-z)}{r}\right)\right],$$

and in this case the policy $(0,0,0)$ can easily be shown to be uniquely optimal (provided $p_1(0,0,0) < 1$). Thus, it is necessary that $\mathbf{a} = (0,0,0)$. However, for $p_0 \in \left(\frac{\alpha}{\lambda\left(\Pi + \frac{CS}{1+\gamma}\right)}, p_1(0,0,0)\right)$, the policy $\mathbf{a}$ cannot induce investment and thus is not optimal. To see this, notice that

$$p_0\lambda\left(\Pi + CS\right) - \alpha - \gamma\alpha\left[\frac{\frac{CS}{1+\gamma}}{\Pi + \frac{CS}{1+\gamma}}\right] > 0 \text{ if } p_0 > \frac{\alpha}{\lambda\left(\Pi + \frac{CS}{1+\gamma}\right)},$$

that is, the expected flow payoff from investment for the government under policy $\left(0, 0, \frac{\frac{CS}{1+\gamma}}{\Pi + \frac{CS}{1+\gamma}}\right)$ is strictly positive; and

$$p_0\lambda\Pi - \left(1 - \frac{\frac{CS}{1+\gamma}}{\Pi + \frac{CS}{1+\gamma}}\right)\alpha > 0 \text{ if } p_0 > \frac{\alpha}{\lambda\left(\Pi + \frac{CS}{1+\gamma}\right)},$$

that is, the expected flow payoff from investment for the firm under policy $\left(0, 0, \frac{\frac{CS}{1+\gamma}}{\Pi + \frac{CS}{1+\gamma}}\right)$ is strictly positive. Therefore the policy $\left(0, 0, \frac{\frac{CS}{1+\gamma}}{\Pi + \frac{CS}{1+\gamma}}\right)$, which can induce investment, is superior to $\mathbf{a}$. This implies that no *ex post* equilibrium exists.■

Proof of Proposition 4:

Consider a pure matching policy $\left(0, 0, \phi^{**}_{\gamma N}\right)$ and the associated abandonment threshold $p_N\left(0, 0, \phi^{**}_{\gamma N}\right) = p^{**}_{\gamma N}$ given by the pair of equations

$$p^{**}_{\gamma N} = \frac{\alpha\left(1 - \phi^{**}_{\gamma N}\right)}{\lambda\Pi}. \tag{25}$$

$$\left\{\lambda p^{**}_{\gamma N}\left[CS + \Pi + (N-1)\,\theta\Pi\right] - \alpha - \gamma\phi^{**}_{\gamma N}\alpha\right\} N k_N\left(p^{**}_{\gamma N}\right) = 0. \tag{26}$$

Conditions (25) and (10) imply

$$\phi_{\gamma N}^{**} = \frac{\frac{CS+(N-1)\theta\Pi}{1+\gamma}}{\frac{CS+(N-1)\theta\Pi}{1+\gamma} + \Pi},$$

$$p_{\gamma N}^{**} = \frac{\alpha}{\lambda\left[\Pi + \frac{CS+(N-1)\theta\Pi}{1+\gamma}\right]}.$$

Now suppose an *ex post* equilibrium exists. Then, in particular, that equilibrium would solve the optimization in (13) for the particular case of $p_0 = 1$. But when $p_0 = 1$, a solution to this optimization is $\mathbf{a} = (0,0,0)$. To see why, substitute $p = 1$ into (12), giving us $W_N(1|\mathbf{a})$. Several steps of algebra reveal

$$W_N(1|\mathbf{a}) = -\frac{\alpha N}{r} + \frac{\lambda N}{(r+\lambda N)}\left\{\frac{r\left[CS + (1+(N-1)\theta)\Pi\right] + N\alpha}{r}\right\} - \frac{N\gamma\left[s + \alpha\phi(1-z)\right]}{(r+\lambda N)}.$$

This expression strictly decreases in s and strictly decreases in ϕ for $z < 1$. It follows that $\max_{\mathbf{a}\in A} W_N(1|\mathbf{a})$ is attained when $\mathbf{a} = (0,0,0)$. Because, by our contrapositive assumption, this solution is an *ex post* equilibrium, the policy $\mathbf{a} = (0,0,0)$ must also solve the optimization in (13) if the government knew that $p_0 \in \left(p_{\gamma N}^{**}, p_N(0,0,0)\right)$. But in such a case the presumptive *ex post* equilibrium policy $(0,0,0)$ will not induce investment and would be inferior to policy $\left(0,0,\phi_{\gamma N}^{**}\right)$, which contradicts $\mathbf{a} = (0,0,0)$ being an *ex post* equilibrium.■

Proof of Proposition 5:

As $W_1(0|\mathbf{a}) \leq 0$ for all $\mathbf{a} \in A$, and $W_1(p|(0,0,0)) \geq 0$ for all $p \in [0,1]$,

$$\max_{\sigma\in\Delta(\Sigma_A)} \min_{\mu\in\Delta(\Sigma)} \int_0^1 \int_A W_1(p|\mathbf{a})\, d\sigma(\mathbf{a})\, d\mu(p) = 0.$$

For any mixed strategy σ assigning positive mass to $\{(z,s,\phi) \in A | s > 0\}$,

$$\min_{\mu\in\Delta(\Sigma)} \int_0^1 \int_A W_1(p|\mathbf{a})\, d\sigma(\mathbf{a})\, d\mu(p) \leq \int_A W_1(0|\mathbf{a})\, d\sigma(\mathbf{a}) < 0,$$

where the first inequality follows upon taking μ to be the distribution that puts full probability mass on $p = 0$, and the second inequality follows from the fact that $W_1(0|\mathbf{a}) < 0$ for any $\mathbf{a} = (z,s,\phi) \in A$ with $s > 0$. Given our requirements that $s \geq \alpha z$ and $z \geq 0$, the support of any mixed strategy max-min solution must thus be contained in $\{(0, 0, \phi)|\phi \in [0, 1]\}$.

Now recall that for $\phi = \phi_{\gamma 1}^{**}$, the socially optimal stopping point is $p_{\gamma 1}^{**}$, giving the government a continuation value of zero. Any $\mathbf{a} = (0, 0, \phi) \in A$ with $\phi > \phi_{\gamma 1}^{**}$ gives rise to a stopping point $p_1(\mathbf{a}) < p_{\gamma 1}^{**}$ and—because of overinvestment and a higher total shadow cost of public funds—continuation values $W_1(p|\mathbf{a}) < 0$ for $p \in (p_1(\mathbf{a}), p_{\gamma 1}^{**}]$. For any mixed strategy σ whose support is contained in $\{(0, 0, \phi)|\phi \in [0, 1]\}$ and which assigns positive probability to policies outside $A_{1\,\max\min}$,

$$\min_{\mu \in \Delta(\Sigma)} \int_0^1 \int_A W_1(p|\mathbf{a})\, d\sigma(\mathbf{a})\, d\mu(p) \leq \int_A W_1\left(p_{\gamma 1}^{**}|\mathbf{a}\right) d\sigma(\mathbf{a}) < 0,$$

where the first inequality follows upon taking μ to be the distribution that puts full probability on $p = p_{\gamma 1}^{**}$, and the second inequality follows from what was shown at the start of this paragraph and the fact that $W_1(p_{\gamma 1}^{**}|0, 0, \phi) = 0$ for $\phi \leq \phi_{\gamma 1}^{**}$.

Finally, since $W_1(0|\mathbf{a}) = 0 \leq W_1(p|\mathbf{a})$ for all $\mathbf{a} \in A_{1\,\max\min}$ and $p \in [0, 1]$, any mixed strategy with support contained in $A_{1\,\max\min}$ satisfies

$$\min_{\mu \in \Delta(\Sigma)} \int_0^1 \int_A W_1(p|\mathbf{a})\, d\sigma(\mathbf{a})\, d\mu(p) = 0$$

and so is the max-min solution.■

References

[1] Almus, M. and D. Czarnitzki (2003), "The Effects of Public R&D Subsidies on Firms Innovation Activities: The Case of Eastern Germany," *The Journal of Business & Economic Statistics*, Vol. 21, No. 2, pp. 226-236.

[2] Arrow, K. J. (1962), "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity,* R. Nelson (ed.) (Princeton, NJ: Princeton University Press), pp. 609-626.

[3] Battigalli, P. and M. Siniscalchi (2003), "Rationalization and Incomplete Information," *Advances in Theoretical Economics*, Article 3.

[4] Bergemann, D. and U. Hege (1998), "Venture Capital Financing, Moral Hazard and Learning," *Journal of Banking and Finance*, Vol. 22, pp. 703-735.

[5] Bergemann, D. and U. Hege (2005), "The Financing of Innovation: Learning and Stopping," *Rand Journal of Economics*, Vol. 36, No. 4, pp. 719-752.

[6] Bergemann, Dirk and Stephen Morris (2007), "Belief Free Incomplete Information Games," Cowles Foundation Discussion Paper No. 1629.

[7] Besanko, D. and J. Wu (2013), "The Impact of Market Structure and Learning on the Tradeoff between R&D Competition and Cooperation," *The Journal of Industrial Economics*, Vol. 61, No. 1, pp. 166–201.

[8] Bonatti, A. and J. Hörner (2011), "Collaborating," *American Economic Review*, Vol. 101, pp. 632–663.

[9] Bodoh-Creed, A. L. (2012), "Ambiguous Beliefs and Mechanism Design," *Games and Economic Behavior*, Vol. 75, No. 2, pp. 518-537.

[10] Bose, S. and E. Ozdenoren and A. Pape (2006), "Optimal Auctions with Ambiguity," *Theoretical Economics*, Vol. 1, No. 4, pp. 411–438.

[11] D'Aspremont, C. and A. Jacquemin (1988), "Cooperative and Noncooperative R&D in Duopoly with Spillovers," *American Economic Review,* Vol. 78, No. 5, pp. 1133-1137.

[12] De Castro, Luciano and Nicholas C. Yannelis (2010), "Ambiguity Aversion Solves the Conflict

Between Efficiency and Incentive Compatibility," working paper.

[13] Gilboa, Itzhak and David Schmeidler (1989), "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, Vol. 18, pp. 141-153.

[14] Hall, B. (2005), "The Financing of Innovation," *Handbook of Technology and Innovation Management*, Scott Shane (ed.) (Blackwell Publishers: Oxford), pp. 409-430.

[15] Hansen, L. P. and T. J. Sargent (2001a), "Acknowledging Misspecification in Macroeconomic Theory," *Review of Economic Dynamics*, Vol. 4, No. 3, pp. 519–35.

[16] Hansen, L. P. and T. J. Sargent (2001b), "Robust Control and Model Uncertainty," *American Economic Review*, Vol. 91, No. 2, pp. 60–6.

[17] Hinloopen, J. (1997), "Subsidizing Cooperative and Noncooperative R&D in Duopoly with Spillovers," *Journal of Economics / Zeitschrift fur Nationalokonomie,* Vol. 66, pp. 151-175.

[18] Hinloopen, J. (2000), "More on Subsidizing Cooperative and Noncooperative R&D in Duopoly with Spillovers," *Journal of Economics / Zeitschrift fur Nationalokonomie,* Vol. 72, pp. 295-308.

[19] Hörner, J., N. Klein, and S. Rady (2014), "Strongly Symmetric Equilibria in Bandit Games," Cowles Foundation Discussion Paper No. 1956.

[20] Hörner, J. and L. Samuelson (2013), "Incentives for Experimenting Agents," *RAND Journal of Economics*, Vol. 44, No. 4, pp. 632–663.

[21] Irwin, D. and P. Klenow (1996), "High-tech R&D Subsidies Estimating the Effects of Sematech," *Journal of International Economics*, Vol. 40, pp. 323-344.

[22] Jones, C. and J. Williams (1998), "Measuring the Social Returns to R&D," *Quarterly Journal of Economics*, Vol. 113, No. 4, pp. 1119-1135.

[23] Kamien, M., E. Muller, and I. Zang (1992), "Research Joint Ventures and R&D Cartels," *The American Economic Review,* Vol. 82, No.5, pp. 1293-1306.

[24] Keller, G., S. Rady, and M. Cripps (2005), "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73, No. 1, pp. 39-68.

[25] Klette, T. J. and Moen, J. (1998), "R&D Investment Responses to R&D Subsidies: A Theoretical Analysis and a Microeconomic Study," University of Oslo working paper.

[26] Klette, T. J. and Moen, J. (1999), "From Growth Theory to Technology Policy: Coordination Problems in Theory and Practice," *Nordic Journal of Political Economy*, Vol. 25, pp. 53-74.

[27] Klein, N. and S. Rady, (2011), "Negatively Correlated Bandits," *Review of Economic Studies*, Vol. 78, pp. 693–732.

[28] Lach, S. (2002), "Do R&D Subsidies Stimulate or Displace Private R&D? Evidence from Israel," *Journal of Industrial Economics*, Vol. 50, No. 4, pp. 369-390.

[29] Lo, Kin Chung (1998), "Sealed Bid Auctions with Uncertainty Averse Bidders," *Economic Theory*, Vol. 12, pp. 1–20.

[30] Pearce, D. (1984), "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, Vol. 52, pp. 279-285.

[31] Qiu, L. D. and Z. Tao, (1998), "Policy on International R&D Cooperation: Subsidy or Tax?," *European Economic Review,* Vol. 42, No. 9, pp. 1727-1750.

[32] Romano, R. E. (1989), "Some Aspects of R&D Subsidization," *Quarterly Journal of Economics*, Vol. 104, No. 4, pp. 863-873.

[33] Rothschild, S. (1974), "A Two-armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, Vol. 9, pp. 185-202.

[34] Sandmo, A. (1998), "Redistribution and the Marginal Cost of Public Funds," *Journal of Public Economics*, Vol. 70, No. 3, pp. 365-382.

[35] Socorro (2007), "Optimal Technology Policy under Asymmetric Information in a Research Joint Venture," *Journal of Economic Behavior & Organization*, Vol. 62, pp. 76-97.

[36] Spencer, B. J. and J. A. Brander (1983), "International R & D Rivalry and Industrial Strategy," *Review of Economic Studies*, Vol. 50, No. 4, pp. 707-722.

[37] Stenbacka, R. and M. Tombak (1998), "Technology Policy and the Organization of R&D," *Journal of Economic Behavior and Organization*, Vol. 36, pp. 503-520.

[38] Takalo, T. and T. Tanayama (2010) "Adverse Selection and Financing of Innovation: Is There a Need for R&D Subsidies?," *Journal of Technology Transfer*, Vol. 35, pp. 16-41.

[39] Wallsten, S. (2000), "The Effects of Government-Industry R&D Programs on Private R&D: The

Case of the Small Business Innovation Research Program," *RAND Journal of Economics*, Vol. 31, No. 1, pp. 82-100.

[40] Xepapadeas, Anastasios (2012), "The Cost of Ambiguity and Robustness in International Pollution Control", in *Climate Change and Common Sense: Essays in Honour of Tom Schelling*, ed. R. W. Hahn, and A. Ulph, Oxford: Oxford University Press.