DISCUSSION PAPER SERIES

DP11939

RESEARCH AND THE APPROVAL PROCESS: THE ORGANIZATION OF PERSUASION

Emeric Henry and Marco Ottaviani

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

RESEARCH AND THE APPROVAL PROCESS: THE ORGANIZATION OF PERSUASION

Emeric Henry and Marco Ottaviani

Discussion Paper DP11939
Published 25 March 2017
Submitted 25 March 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

RESEARCH AND THE APPROVAL PROCESS: THE ORGANIZATION OF PERSUASION

Abstract

An informer sequentially collects information at a cost to influence an evaluator's choice between rejection and approval. Payoffs and control rights are split between informer and evaluator depending on the organizational rules governing the approval process. We compare the performance of different organizations from a positive and normative perspective, depending on the commitment power of informer and evaluator. As a welfare benchmark we recover Wald's (1947) classic solution for a statistician with payoff equal to the sum of our informer and evaluator. We apply the analysis to the regulatory process for drug approval and to the market for new technologies.

JEL Classification: D83, M38

Keywords: Persuasion, Information, Organization, Opproval

Emeric Henry - emeric.henry@sciences-po.fr
Sciences Po and CEPR

Marco Ottaviani - marco.ottaviani@unibocconi.it
Bocconi University and CEPR

Acknowledgements
Ottaviani acknowledges financial support from the European Research Council through ERC Advanced Grant 295835. We thank Alessandro Bonatti, Marina Halac, and Joe Harrington for helpful discussion, Luca Ferocino and Paolo Varraso for excellent research assistance, and seminar audiences at Amsterdam, Barcelona, Bergamo, Bordeaux, CEPR 2014 IMO Workshop in Frankfurt, Capri, Columbia, Copenhagen, EARIE 2013 in Evora, Lausanne, Glasgow, ESSET 2014 and 2015 in Gerzensee, ESEM 2015 in Mannheim, MIT, Oslo, Paris, Petralia 2014, Warwick, and Zurich.

Research and the Approval Process: The Organization of Persuasion*

Emeric Henry† Marco Ottaviani‡

March 2017

Abstract

An informer sequentially collects information at a cost to influence an evaluator's choice between rejection and approval. Payoffs and control rights are split between informer and evaluator depending on the organizational rules governing the approval process. We compare the performance of different organizations from a positive and normative perspective, depending on the commitment power of informer and evaluator. As a welfare benchmark we recover Wald's (1947) classic solution for a statistician with payoff equal to the sum of our informer and evaluator. We apply the analysis to the regulatory process for drug approval and to the market for new technologies.

Keywords: Persuasion, information, organization, approval.

JEL Classification: D83 (Search; Learning; Information and Knowledge; Communication; Belief), M38 (Government Policy and Regulation).

*Ottaviani acknowledges financial support from the European Research Council through ERC Advanced Grant 295835. We thank Alessandro Bonatti, Marina Halac, and Joe Harrington for helpful discussion, Luca Ferocino and Paolo Varraso for excellent research assistance, and seminar audiences at Amsterdam, Barcelona, Bergamo, Bordeaux, CEPR 2014 IMO Workshop in Frankfurt, Capri, Columbia, Copenhagen, EARIE 2013 in Evora, Lausanne, Glasgow, ESSET 2014 and 2015 in Gerzensee, ESEM 2015 in Mannheim, MIT, Oslo, Paris, Petralia 2014, Warwick, and Zurich.

†Sciences Po and CEPR.

‡Bocconi and CEPR.

1 Introduction

To persuade doctors or a regulatory authority to approve a drug, a pharmaceutical company performs costly clinical trials that document the drug's effectiveness and safety. Similarly, a developer of a new technology can test it to convince potential acquirers to buy. Or a company's division can collect evidence on a new product's profitability to get headquarters on board for launching the product.

In these situations, a biased informer attempts to persuade an evaluator through costly information diffusion. The organization of the interaction between informer and evaluator varies across settings. In some instances the procedure followed involves limited commitment on the side of both players. For example, consider an author who can conduct additional tests to convince an editor to accept a paper. The author can submit at any point in time and the editor can always ask for more information until final acceptance. The editor does not typically commit to a standard for acceptance, and the author does not commit not to do more research if the editor were to ask for it. Instead, in the context of drug approval, regulators such as the US Food and Drug Administration (FDA) can be seen as currently committing to an approval standard before research starts, for example by defining the margin of error.[1] This paper develops a flexible framework to analyze and welfare rank different institutions that govern persuasion, depending on the extensive-form game induced by the organizational structure.

Our approach is based on a strategic deconstruction of Abraham Wald's (1947) model of sequential information acquisition. Wald features a single player, a statistician who not only decides on—and pays for—information acquisition but also controls the final approval/rejection decision. From a technical perspective, we obtain analytical tractability by formulating the problem in continuous time.[2] Information collection is modeled as a stochastic process whose drift depends on a binary state, either good or bad, corresponding to whether the statistician prefers approval or rejection. The outcome of an infinitesimal experiment is observed in each instant in which research is conducted at cost c. Wald's decision-theoretic solution to the statistician's sequential information acquisition problem with exponential discounting at rate r is characterized by an (s,S) policy with two standards. The statistician optimally conducts research when the posterior belief about the good state based on the information acquired in the past lies between the two standards (s,S). The

[1] See Section 4 on the historical evolution of pharmaceutical regulation.

[2] Our continuous-time formulation builds on the decision-theoretic treatments by Dvoretsky, Keifer, and Wolfowitz (1953), Mikhalevich (1958), and Shiryaev (1967); see Shiryaev (1978, Chapter IV, Sections 1-2) for a textbook treatment. See also Moscarini and Smith (2001) for a characterization of the intensive margin of sequential experimentation in a non-strategic environment; our formulation focuses on the simpler case with one experiment per period.

statistician either approves as soon as the belief becomes sufficiently favorable and hits the upper approval standard S or rejects as soon as the belief becomes sufficiently unfavorable and hits the lower rejection standard s.

While in Wald's classic framework a single statistician controls both research and final approval/rejection, our strategic deconstruction splits control of research and final decision between two players, an informer i and an evaluator e. Rejection yields a zero payoff to the informer as well as to the evaluator, who both discount future payoffs at the same rate. Approval gives the evaluator a positive benefit $v_e^G > 0$ if the state is good and a negative benefit $v_e^B < 0$ if the state is bad. The informer bears the cost of research and obtains a fixed benefit $v_i > 0$ from approval regardless of the state. Wald's statistician w can be seen as a reconstructed social planner who maximizes the sum of the payoffs of informer and evaluator.

The first contribution of the paper is the equilibrium characterization for the Wald persuasion games we introduce in Section 2. To set the stage, Section 2.1 considers the *informer-authority* game in which the informer can make a single take-it-or-leave-it approval request to the evaluator, thereby committing to stop research in case of rejection. Under informer authority, the informer's research strategy follows an (s,S) policy. The upper standard, at which the informer stops research and submits the approval request is the evaluator's myopic cutoff, equal to the belief at which the evaluator, when required to make an immediate decision, is indifferent between rejection and approval. For intermediate beliefs the informer continues conducting research, trading off the expected cost of research against the benefit from potential approval if the myopic cutoff is reached. If the belief becomes sufficiently pessimistic, the informer abandons research because the expected cost of additional information is higher than the marginal benefit of persuasion. Note that the informer has no interest in information per se, given that the informer's payoff from approval does not depend on the underlying state. Nevertheless, the informer values information instrumentally in order to persuade the evaluator to approve.

Informer authority can be seen as a continuous-time limit of Brocas and Carrillo's (2007) discrete-time model, with the added generality of allowing also for payoff discounting as well as costly information. The structure of the solution is closely related to the one characterized by Kamenica and Gentzkow (2011), KG from now on, for the case in which the informer (sender) can choose the optimal information structure without any constraint other than Bayesian rationality by the evaluator (receiver). Compared to KG, our informer is unable to commit to the information structure and is restricted to choose at each instant whether or not to obtain a signal generated by Brownian diffusion. In the limit, as research cost and discount rate both go to zero, the outcome of our dynamic game with informer authority converges to KG's unconstrained solution. Thus, KG's assumption of commitment to the signal structure can be dispensed with, given that the solution of the informer's sequential Wald problem is dynamically consistent.

This connection between Wald and KG is the point of departure for our analysis. By allowing for costly and sequential acquisition of information à la Wald, we embed persuasion in a game-theoretic framework that can be applied to clinical trials and other settings with strategic information diffusion. The organization of the approval process determines the extensive-form bargaining protocol that governs the players' interaction. As we show, a number of dynamic games capturing natural organizational structures can be reduced to corresponding static games amenable to simple analysis.

Section 2.2 considers the *no-commitment* case in which neither the informer nor the evaluator can commit to a policy. Compared to informer authority, the evaluator is now more powerful and is able to request further research when not satisfied with the current level. The resulting Nash equilibrium is at the intersection of the informer's lower best reply (optimal choice of the rejection standard for a given approval standard) and the evaluator's upper best reply (optimal choice of the approval standard for a given rejection standard). In this no-commitment outcome, both the approval standard and the rejection standard are higher than under informer authority. The evaluator values information and does not pay its direct cost but suffers from delay in obtaining the approval payoff. This leads the evaluator to always set a standard higher than the myopic cutoff. In turn, the rejection standard controlled by the informer increases compared to the informer-authority solution. Intuitively, the informer is discouraged from acquiring information at the bottom, given that it is harder to reach the higher approval standard set by the evaluator in equilibrium.

Section 2.3 turns to *evaluator commitment*. When the evaluator can commit at the outset by setting a standard that the evidence must surpass for approval to be granted, the evaluator chooses the preferred point on the informer's lower best reply. In the resulting Stackelberg outcome, we show that the evaluator commits to an approval standard that is necessarily below the Nash level. To see why, note that the informer undertakes too little research at the bottom in the eye of the evaluator because the informer must pay the cost of research while the evaluator free rides on the information. Improving the informer's incentives to undertake research at the bottom results in a first-order gain for the evaluator. It is then optimal for the evaluator to set a more lenient (i.e., lower) approval standard than in the Nash outcome, so as to reduce the informer's rejection standard and thus induce more information collection.

Wald persuasion games connect persuasion with the literature on experimentation. Seen from the perspective of the experimentation literature, Wald persuasion games incorporate into Wald's (1947) sequential analysis the strategic issues that arise when information collection and final decision are made by two different players. Relative to work on persuasion, our first scenario with informer authority corresponds to a continuous-time version of Brocas and Carrillo (2007) with discounting and costly research, obtaining KG as the frictionless limit. The continuous-time formulation we adopt is close to Gul and Pesendorfer (2012) and Chan, Lizzeri, Suen, and Yariv (2015) who consider strategic settings in which public information arrives over time to voters.[3] While in their setting information is revealed publicly to all voters, we focus on the sequential interaction between an informer who collects information and then reports it to an evaluator who makes the approval decision. To the persuasion literature we contribute the characterization of the equilibrium in Wald persuasion games with different commitment structures.

In a recent contribution also on the interplay between agency and experimentation, Guo (2016) analyze how a principal should dynamically delegate experimentation to a privately-informed but biased agent in an exponential bandit model à la Keller, Rady, and Cripps (2005).[4] Instead, we assume away pre-existing private information by the informer, but relax the full commitment assumption in a variety of ways. In their analysis of investment under uncertainty with information held by a biased agent, Grenadier, Malenko, and Malenko (2016) compare solutions under full commitment and no commitment. Our tack, instead, is to consider intermediate scenarios for commitment; to this end, we allow the informer to disclose the current state of information when applying for approval.[5]

Our second contribution is Section 3's welfare comparison of informer authority, no commitment, and evaluator commitment. Given that the evaluator finds it optimal to commit to be more lenient than under Nash, commitment benefits not only the evaluator but also the informer. Thus, evaluator commitment always Pareto dominates no commitment. The comparison between informer authority and evaluator commitment is more subtle.

With costless research and no discounting, KG's informer authority leads to zero false negatives

[3]In Gul and Pesendorfer's (2012) model, information is provided by the party that leads, whereas in Chan, Lizzeri, Suen, and Yariv (2015) voters decide collectively themselves when to stop acquiring public information and reach a decision.

[4]More broadly, this paper is related to the social experimentation literature spearheaded by Bolton and Harris (1999). While that literature focuses on incentives for multiple experimenters in bandit models, we focus on the interaction between a single experimenter and a decision maker.

[5]In concurrent work, Orlov, Skrzypacz, and Zryumov (2016) and Bizzotto, Jesper and Vigier (2016) consider instead dynamic persuasion games in which information flows exogenously and the agent at each instant has the power to design any information structure as in KG. In those papers, the receiver can exercise the option at any time without need for submission by the sender—another key difference from the extensive-form games we consider.

(type II errors), but to a socially excessive amount of false positives (type I errors). In the context of drug approval, in this frictionless benchmark the pharma industry's ability to control information results in the approval of too many harmful drugs. Evaluator commitment, instead, results in full information and thus attains the planner's first-best outcome, with no error of either type.

However, this frictionless welfare ranking is largely overturned once we introduce discounting and/or information costs, as we show. Frictions tend to make informer authority socially preferred to evaluator commitment, through two channels. First, when research is costly, all organizational forms always lead to excessive false negatives. The informer abandons research inefficiently too early because the evaluator retains some veto power and free rides on the information.[6] Second, with sufficient discounting or costly research, false positives become socially insufficient—rather than excessive—because the evaluator does not internalize the informer's benefits from adoption and thus approves too little. Giving authority to the informer alleviates both problems. In the context of drug approval, transferring authority to associations of doctors when research is costly leads to insufficient incentives for pharmaceutical companies to conduct clinical trials and to limited desire of doctors to adopt new drugs.

For our third contribution, we put the framework to work in two applied contexts:

- Section 4 zooms in on our motivating application to approval regulation and discusses how the drug approval process evolved in the US. Before any regulation the organization of drug approval in the US resembled informer authority, followed by no commitment from 1905 when the Council of Pharmacy and Chemistry was formed within the American Medical Association. After the creation of the FDA in 1927, the role of the evaluator was taken up by the social planner, initially with no commitment and with social planner commitment following 1962's reform. As we show, for low-price drugs it is optimal for the planner to delegate power to the evaluator under informer authority, but for high-price drugs the planner prefers to take up the role of player. Under planner commitment, the regulator should commit to be more lenient to encourage more research for orphan drugs with relatively low market value. For high-price drugs that pharmaceutical companies find attractive to commercialize, instead, the regulator should commit to a more exacting standard of approval than under no commitment to discourage excessive testing.

[6] In our setting, excessive rejection stems from the fact that the evaluator free rides on the informer's costly effort to learn about quality. In Green and Taylor (2016), instead, rejection is socially excessive because it serves as an incentive device, inducing the agent to report truthfully the progress made on the project.

- Section 5 extends the model to fit the market for technology transfer by a developer selling to competing downstream producers. This application allows us to introduce endogenous pricing. We show that if producers exert negative externalities on their competitors when adopting a new technology, the organization of persuasion matters. In particular we show that producers are made worse off by the introduction of a new technology. This setting also provides a unifying framework incorporating Gilbert and Newbery's (1982) preemptive bidding effect and Reinganum's (1983) replacement effect in the special case where the first evaluator is an incumbent and the second a potential entrant.

2 Wald Persuasion Games

Two players, an informer i and an evaluator e, interact in continuous time under uncertainty about the state of the world ω, which can be either good G or bad B. The decision to be made, either approval A or rejection R, is irreversible. The payoff from rejection is zero for all players, regardless of the state. The evaluator's payoff from approval is positive $v_e^G > 0$ in the good state but negative $v_e^B < 0$ in the bad state; the informer obtains a state-independent benefit from approval equal to $v_i > 0$.[7]

At the outset $t = 0$ players share a common prior about the state $q_0 = \Pr(\omega = G)$. We denote $\hat{q}_e$ the myopic cutoff corresponding to the belief at which the evaluator is indifferent between the two alternatives A and R,

$$\hat{q}_e\, v_e^G \;+\; (1-\hat{q}_e)\, v_e^B = 0.$$

If forced to make a decision at belief q, the evaluator chooses A if and only if $q \geq \hat{q}_e$. Clearly, the informer's myopic cutoff is $\hat{q}_i = 0$.

At each instant the informer can conduct research whose results are publicly disseminated. The arrival of new information is modeled as a Wiener process $d\Sigma$, whose drift is determined by the state. Specifically, the process has positive drift μ and variance ρ^2 if the state is G or drift $-\mu$ and variance ρ^2 if the state is B. Accumulating information over a period of time dt costs cdt. Finally, both players discount future payoffs at the same rate $r \geq 0$. The informer pays the cost of research, while the evaluator free rides on the information publicly revealed.

Wald Benchmark. As a welfare benchmark for comparison, consider the problem of a social planner playing the role of Wald's (1947) statistician who controls both research and approval decisions, and obtains the sum of the informer's payoff and the evaluator's payoff, including the cost

[7] The model can be extended to allow for a more general state-dependent structure; see the discussion in Section 6.

of research. Thus, approval in state ω results in decision payoff $v_w^\omega = v_i + v_e^\omega$ for the planner. The solution is well known:

Proposition 0 *The Wald solution consists of two standards s_w^* (the rejection standard) and S_w^* (the approval standard), such that the planner:*

- *stops researching and rejects if $q \leq s_w^*$,*
- *conducts research if $s_w^* < q < S_w^*$, and*
- *stops researching and approves if $q \geq S_w^*$.*

See Supplementary Appendix B for a self-contained proof. The problem is solved by finding the simultaneous solution of the two first-order (also known as smooth-pasting) conditions determining (1) the choice of the rejection standard $s = b_w(S)$ for a given approval standard S and (2) the choice of the approval standard $S = B_w(s)$ for a given s.

Wald Deconstructed. Our model is a deconstruction of Wald with (a) research costs borne by the informer, (b) payoff from approval split between the informer and the evaluator, and (c) authority over research and approval decisions allocated to the players depending on the organizational structure.

The derivation of the equilibria resulting in the different organizations we consider builds on the best replies of informer and evaluator. To derive the informer-authority solution, Section 2.1 characterizes the informer's choice of the lower standard $s = b_i(S)$ as a best reply to a given upper standard S. Section 2.2 then characterizes the evaluator's choice of $S = B_e(s)$ as a best reply to s and derives the no-commitment solution, which relies on both the informer's lower best reply and the evaluator's upper best reply. Finally, Section 2.3 analyzes the evaluator-commitment solution, corresponding to the Stackelberg outcome in which the evaluator chooses the most preferred S on the informer's best reply $s = b_i(S)$. We return to the Wald solution in Section 3 when comparing the welfare achieved by the different organizations we consider.

2.1 Informer Authority

In the first organizational form we consider, the informer performs research and can make one take-it-or-leave-it demand for approval to the evaluator, who must then choose once and for all between rejection and approval. This can be seen as a situation where the informer commits not to research further in case of rejection or, alternatively, as a situation where the evaluator has no power to mandate further research.

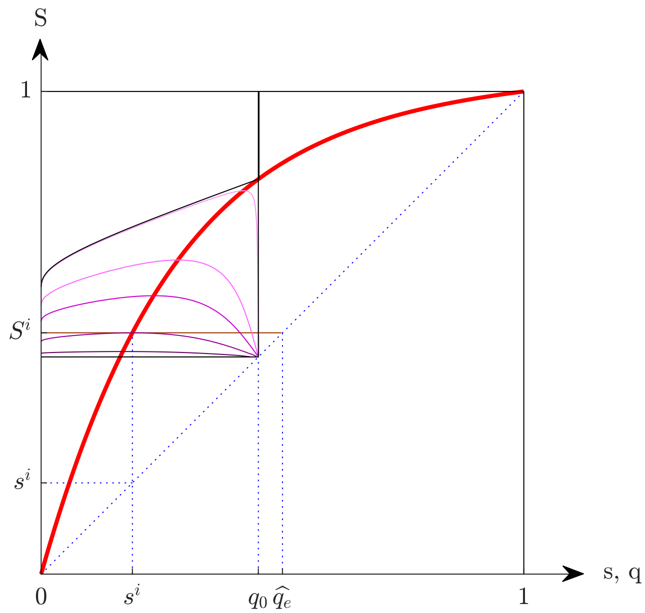

Figure 1: Informer's best reply $s = b_i(S)$.

Specifically, we study the stationary Markov Perfect Equilibrium (MPE) of the following dynamic game with informer authority. In each instant t, the evaluator and the informer move sequentially according to the following rules:

1. The informer chooses between $\mathscr{I}_i$ (information acquisition), $\mathscr{A}_i$ (application for approval), or $\mathscr{R}_i$ (rejection);

2. If the informer chooses $\mathscr{A}_i$, then the evaluator chooses between $\mathscr{A}_e$ (approval) or $\mathscr{R}_e$ (rejection);

3. The outcome of the game is I (information acquisition) if $(\mathscr{I}_i)$, A (approval) if $(\mathscr{A}_i, \mathscr{A}_e)$, and R (rejection) if either $(\mathscr{R}_i)$ or $(\mathscr{A}_i, \mathscr{R}_e)$.

The state variable is represented by the belief q.

The informer does not care about information per se but just uses it to convince the evaluator. The informer thus stops research as soon as there is sufficient evidence to induce the evaluator to approve, i.e. as soon as the belief q reaches the myopic cutoff $\hat{q}_e$.

To determine the research decision of the informer for values of q below $\hat{q}_e$ it is necessary to derive the informer's lower best reply $s = b_i(S)$, plotted as the thick line in Figure 1. To visualize the construction, fix an approval standard $S \in [0,1]$ on the vertical axis and consider a prior belief $q_0 < S$ on the horizontal axis, which tracks both q_0 and $s \in [0,1]$. The prior belief q_0 defines a quadrant of attainable (s,S) values and in this space we plot the iso-payoff curves of the informer. The informer's bliss point is at the origin, $(s_i^* = 0, S_i^* = 0)$, since the informer with full control on the decision would approve immediately. Thus, iso-payoff curves closer to the informer's bliss point correspond to higher payoff levels. To a given approval standard S, the informer best replies by choosing the rejection standard s that lies on the highest iso-payoff curve attainable at S. The figure illustrates this choice satisfying tangency to the horizontal line $S = S^i$.

When choosing the best-reply rejection standard, $s = b_i(S)$, the informer trades off the expected discounted cost of research needed to reach S against the expected gain from undertaking research, corresponding to the expected discounted probability that S will be reached, which will result in the accrual of v_i,

$$\underbrace{\beta_1(s,S)\, v_i}_{\text{expected gain from reaching } S} = \underbrace{\beta_2(s,S)\, c/r}_{\text{financial cost of research}} \tag{1}$$

where $\beta_1(s,S) > 0$ and $\beta_2(s,S) > 0$ are defined in Appendix A. The following results hold:

Proposition 1 ***(a)*** *The informer's lower best reply $b_i(S)$ is* ***(i)*** *independent of the current belief q and* ***(ii)*** *increasing in S.*

(b) *The unique MPE outcome of the informer-authority game is a Wald-cutoff path with standards (s^i, S^i) such that $S^i = \hat{q}_e$ and $s^i = b_i(\hat{q}_e)$.*

(c) *The equilibrium value of the informer at the prior belief q_0 converges to the value of the optimal signal characterized by KG when both the discount factor r and the cost of research c converge to zero.*

The comparative statics in part (a) holds because a higher S, corresponding to a request for more information at the top before approval, increases the research cost and moves the informer away from the bliss point $(s_i^* = 0, S_i^* = 0)$. It is then optimal for the informer to abandon research earlier, for a higher s, corresponding to less information at the bottom. Thus, the need for *more* information at the top induces the informer to provide *less* information at the bottom. In this sense, the informer perceives information provision to be a strategic substitute.[8]

For part (b), the evaluator approves as soon as the belief reaches the myopic cutoff $\hat{q}_e$, i.e., for any $q \geq \hat{q}_e$. At the lower end the informer abandons research as soon as the belief q falls below the lower best reply to $\hat{q}_e$, resulting in $\left(s^i = b\left(\hat{q}_e\right), S^i = \hat{q}_e\right)$, as represented in Figure 1.

Part (c) turns to the comparison to KG:

[8]See Brocas, Carrillo, and Palfrey (2012) for experimental evidence of strategic substitutability along these lines.

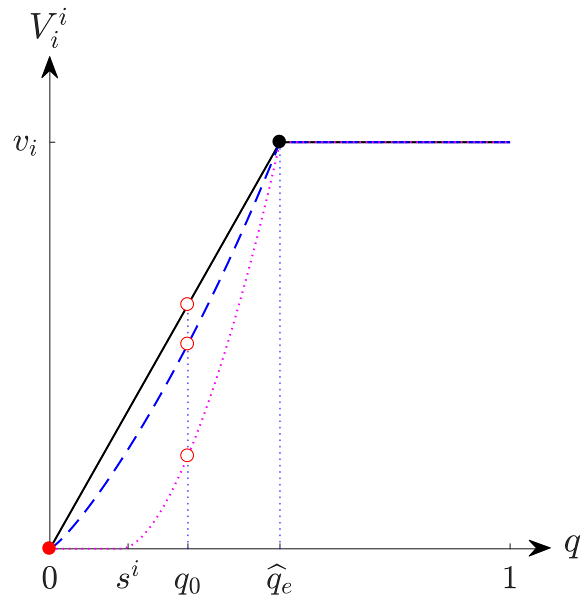

Figure 2: Informer's expected value under informer authority against the belief, for $c = 0, r = 0$ (continuous curve), $c = 0, r = 3$ (dashed curve), and $c = 4, r = 3$ (dotted curve).

- KG allow the informer to commit to any signal structure. They show that the optimal signal structure has the following two properties. First, when the evaluator takes the informer's preferred action, the evaluator is exactly indifferent between approval and rejection. Second, when taking the informer's least preferred decision, the evaluator is completely certain of the state. Thus, the informer achieves the optimal solution through an "extremal" binary signal with an asymmetric conditional distribution taking the prior q_0 to posterior $\hat{q}_e$ with probability $q_0/\hat{q}_e$ and to posterior 0 with complementary probability. Figure 2, which echoes KG's Figure 2, plots the informer's value function under informer authority, denoted by $V_i^i(q)$, against the belief q[9] The informer's value obtained by KG through concavification is equal to the continuous curve, given that the evaluator rejects whenever the belief is below $\hat{q}_e$ (giving a payoff of zero to the informer) and approves above $\hat{q}_e$ (yielding v_i).

- Instead, we only allow our informer to choose a signal in a particular class of Brownian diffusion signals, without commitment. However, we find that this restriction is inessential. First, at the top the informer immediately stops researching when the belief reaches the evaluator's myopic cutoff, exactly as in KG. In analogy with KG's second property, the signal induced in our equilibrium is also extremal. Discounting of payoffs convexifies the concavified portion of the value function where learning takes place, thus reducing the informer's payoff, as illus-

[9] With our notation, KG's optimal signal has $\Pr(S|B) = 1 - \Pr(s|B) = \frac{q_0}{1-q_0}\frac{1-\hat{q}_e}{\hat{q}_e}, \Pr(S|G) = 1 - \Pr(s|G) = 1$.

trated by the dashed value function in Figure 2. When the informer takes into account the cost of information collection, $c > 0$, the informer abandons research when the belief is sufficiently unfavorable and hits the rejection standard $s = b_i(\hat{q}_e) > 0$, where the dotted curve is equal to zero. Thus, information costs reduce the informer's incentives for information collection.

When the research cost c vanishes, the belief $s^i = b_i(\hat{q}_e)$ at which the informer stops researching and induces rejection converges to 0 and the informer's equilibrium value function in the game with informer authority is equal to the one characterized by KG. More generally, time consistency of the solution of the sequential Wald problem means that KG's assumption of commitment to the signal structure can be dispensed within our dynamic implementation. As we see more generally in the remainder of the paper, the equilibrium signal structure is extremal when informer and evaluator interact in a number of other realistic ways.

2.2 No Commitment

Our Wald approach allows us to analyze the equilibrium amount of persuasion resulting from a variety of extensive-form games that characterize the interaction of informer and evaluator, other than the informer-authority case which has been the focus of most of the literature on persuasion. Consider now an organization with a more powerful evaluator that can refrain from making a decision if not satisfied with the evidence presented.[10] Specifically, modify the informer-authority game described in the previous section in the following way: (i) at stage 2, following application for approval by the informer ($\mathscr{A}_i$), allow the evaluator to also choose request further information $\mathscr{I}_e$ (information), in addition to $\mathscr{A}_e$ (approval) or $\mathscr{R}_e$ (rejection); and (ii) at stage 3, specify that following $(\mathscr{A}_i, \mathscr{I}_e)$ the outcome of the game is I (information acquisition).

The equilibrium relies on the characterization of the evaluator's upper best reply. Note two differences between the evaluator's problem and the informer's problem analyzed in Proposition 1.a. First, the evaluator does not bear the cost of research. Second, the evaluator cares about the information revealed in the research process. Even though research has no financial cost for the evaluator, discounting induces an opportunity cost of delaying a good decision.

For a given lower threshold s, the evaluator optimally chooses S so as to tradeoff the value of information and the cost of delaying the decision. For $s < \hat{q}_e$, the first-order condition characterizing

[10] Alternatively, the informer might be unable to commit not to research further in case of rejection.

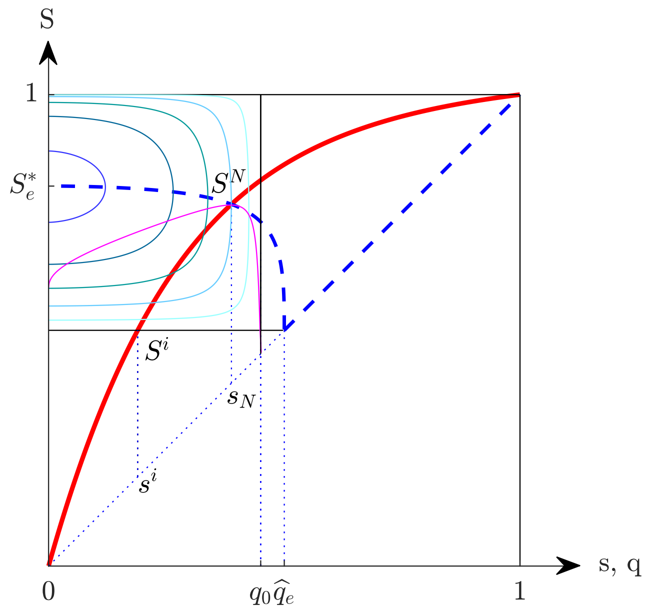

Figure 3: No-commitment solution.

the evaluator's best reply $S = B_e(s)$ reflects this tradeoff

$$\underbrace{-\beta_3(s,S)\, v_e^B}_{\text{benefit of information}} = \underbrace{\beta_4(s,S)\, v_A(S),}_{\text{cost/ benefit of delaying decision}} \tag{2}$$

where $\beta_3(s,S) > 0$ and $\beta_4(s,S)$ are defined in Appendix A and

$$v_A(S) = S\, v_e^G + (1-S)\, v_e^B \tag{3}$$

is the expected benefit of approval at belief $q = S$, which is positive for $q > \hat{q}_e$.

Figure 3 shows $B_e(s)$ in the segment $s < \hat{q}_e$ as the downward sloping dashed curve. As in the case of the lower best reply, to visualize the construction, fix a rejection standard $s \in [0,1]$ on the horizontal axis and consider a prior belief $q_0 < S$ on horizontal axis, which tracks both q_0 and $s \in [0,1]$. For a given prior belief q_0 we plot the iso-payoff curves of the evaluator. Given that the evaluator does not pay for information, the evaluator's bliss point features $s_e^* = 0$, but $S_e^* \in (\hat{q}_e, 1)$ because delaying a good decision has a positive opportunity cost due to discounting. Iso-payoff curves closer to the evaluator's bliss point correspond to higher payoff levels. The evaluator chooses S on the highest iso-payoff curve crossing the given approval standard s, at the point of tangency with the vertical line at s. The following results hold:

Proposition 2 *(a) The evaluator's upper best reply $B_e(s)$ is **(i)** independent of the current belief q and **(ii)** decreasing in s for $s < \hat{q}_e$ and equal to s for $s \geq \hat{q}_e$.*

(b) *The unique MPE outcome of the no-commitment game is a Wald-cutoff path with standards* (s^N, S^N) *such that* $s^N = b_i(S^N)$ *and* $S^N = B_e(s^N)$*.*

(c) *Compared to the informer-authority solution, more information is obtained at the top* $S^N > S^i$ *and less at the bottom* $s^N > s^i$*.*

As claimed in part (a), for $s \geq \hat{q}_e$, it is optimal for the evaluator to immediately approve because information has a negative value, thus $B_e(s) = s$ along the diagonal.[11] To see why $B_e(s)$ is decreasing in s for $s < \hat{q}_e$, note that any $s > s_e^* = 0$ results in too little information at the bottom for the evaluator. A higher s, corresponding to less information at the bottom, reduces the evaluator's marginal value of information at the top. The evaluator then best replies by approving earlier for a lower S to avoid delaying the benefits $v_A(S)$. Thus, the acquisition of *less* information at the bottom induces the evaluator to require *less* information at the top. In this sense, the evaluator perceives information provision to be a strategic complement, rather than substitute as the informer.

For part (b), the unique stationary MPE outcome of this game (s^N, S^N) is displayed in Figure 4 at the intersection of the evaluator's upper best reply $S^N = B_e(s^N)$ with the informer's lower best reply $s^N = b_i(S^N)$, described in the previous section. This is the same outcome as the Nash equilibrium of a one-shot simultaneous-move game in which the informer chooses the lower standard and the evaluator chooses the upper standard.

For part (c), recall that the evaluator values information but does not pay the direct cost of research. The evaluator only takes into account the opportunity cost of delaying approval with positive expected value. At the myopic cutoff the evaluator's expected value of approval is zero, so delaying has no opportunity cost. Thus, it is optimal for the evaluator under no commitment to delay approval beyond the myopic cutoff $\hat{q}_e$; the Nash approval standard satisfies $S^N > S^i = \hat{q}_e$. This delay in turn results in $s^N > s^i$. Given that the approval standard is set at a level further away from the unconstrained optimal level for the informer, $\hat{q}_i = 0$, the value of research at the bottom for the informer is reduced by Proposition 1.a.ii.

2.3 Evaluator Commitment

The evaluator is weakly better off under no commitment than under informer commitment, and strictly better off for a prior in (s^i, S^N). The evaluator can do even better by initially committing to

[11] To see this, suppose that $s \geq \hat{q}_e$ and consider a belief $q \geq s$. If the evaluator were to set $S > q$, with probability q/s the belief would reach s (resulting in rejection and thus in payoff 0 for the evaluator) and with probability q/S the belief would reach S (with approval payoff $v_A(S)$). Using (3) the resulting expected payoff for the evaluator is clearly below $qv_e^G + (1-q)v_e^B$, which can be obtained through immediate approval. Thus, it is optimal to set $S = q$ for all $q \geq s$.

approve according to an ex-ante specified rule. Specifically, suppose that at the start of the game the evaluator can commit to an approval rule that depends only on the belief at the time of decision, but not on the path or time taken to get there. As we show, the optimal approval rule takes the following cutoff form: approve if and only if $q \geq S^e$.

Under evaluator commitment, the evaluator chooses the preferred point on the informer's lower best reply. We show in the following result that for a fixed initial belief at an intermediate level, $q_0 \in (s^i, S^N)$, it is optimal for the evaluator to commit to an approval standard below the no-commitment level S^N in order to encourage the informer to perform more research at the bottom. The tradeoff is between:

- *A second-order negative direct effect:* Holding fixed the informer's strategy s, decreasing the upper standard below the evaluator's upper best reply $S = B_e(s)$, means that an insufficient amount of research is performed at the upper end. This loss is clearly second order by the envelope theorem because we start from the evaluator's optimal choice of S holding fixed the informer's choice of s.

- *A first-order positive strategic effect:* The informer's strategic reply to the decrease in the commitment S is to increase research at the lower end. Given that the informer's initial choice of s was higher than what the evaluator would have liked, this is a positive first-order effect for the evaluator that dominates the second-order negative effect.

The ability to commit strictly benefits the evaluator for $q_0 \in (s^i, S^N)$, i.e., whenever the optimal commitment is interior:

Proposition 3 *In the evaluator-commitment game:* ***(i)*** *if $q_0 \in (s^i, S^N)$ the evaluator chooses an interior commitment $S^e(q_0)$ increasing in q_0 and such that $S^e(q_0) \in (\hat{q}_e, S^N)$;* ***(ii)*** *if $q_0 < s^i$ no research is performed regardless of the choice of the evaluator; and* ***(iii)*** *if $q_0 > S^N$ the evaluator chooses to immediately approve.*

A key property of the commitment path is that dynamic consistency no longer applies. The optimal choice of commitment by the evaluator depends on the initial belief q_0 at the time of commitment. As the prior belief becomes more favorable, the strategic effect, capturing the benefit of encouraging the informer, becomes less relevant since it becomes very likely that the approval standard will be quickly reached. This explains why the interior commitment is increasing in q_0.

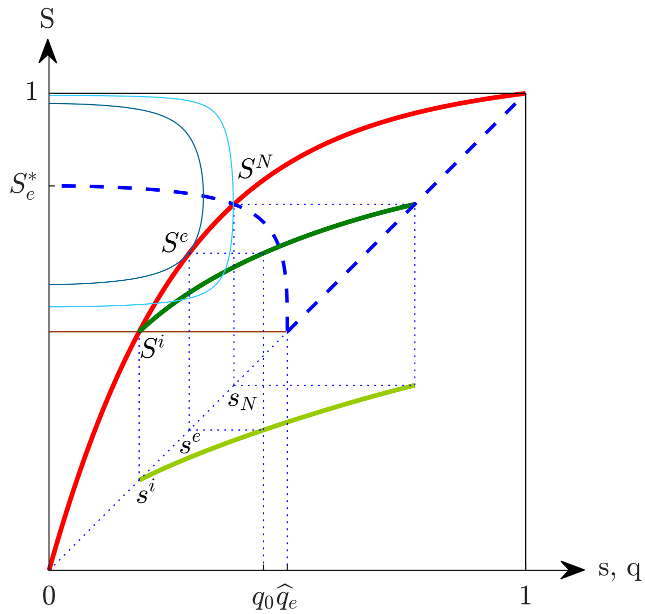

Figure 4: Best replies and commitment path.

Figure 4 displays the commitment path $S^e(q_0)$ as the upper bold curve plotted in green, as a function of the initial belief q_0 rather than s on the horizontal axis; the lower bold curve plotted in light green corresponds to the resulting $s^e(q_0) = b_e(S^e(q_0))$. To understand the construction, fix an initial belief $q_0 \in (s^i, S^N)$ and plot the iso-payoff curves of the evaluator. The optimal commitment is obtained as the tangency point between the highest iso-payoff curve of the evaluator and the lower best reply of the informer $b_i(S)$; the iso-payoff curves to the left, corresponding to higher levels of payoff for the evaluator, are not attainable. The commitment path is obtained by repeating this procedure for each value of q_0. The optimal commitment is above the myopic cutoff $\hat{q}_e$, because otherwise the evaluator would obtain a negative payoff. When the initial belief is below s^i, any commitment above $\hat{q}_e$ will push the informer to do no research, so that in this region commitment has no value for the evaluator.[12] At the other extreme, when the prior belief is above S^N, it is optimal for the evaluator to choose immediate approval which induces the informer to do no research.[13]

[12]By definition of s^i, the expected cost of research dominates the expected gain, even in the best-case scenario from the informer's point of view where the standard would be set exactly at the myopic cutoff.

[13]Note that there are more initial beliefs q_0 for which some research is conducted under evaluator commitment than under no commitment: $(s^N, S^N) \subset (s^i, S^N)$.

3 Welfare Comparison of Organizations: Wald Reconstructed

This section compares the welfare performance of the three decentralized organizations introduced so far: informer authority, no commitment and evaluator commitment. The normative benchmark is the social planner's Wald solution (s_w^*, S_w^*) that results from a centralized setting in which the planner controls both the research and the approval decisions, and maximizes the sum of the informer's and the evaluator's payoffs, taking into account the cost of research. The Wald solution (s_w^*, S_w^*), characterized in Proposition 0 as the intersection between the planner's optimality conditions $s = b_w(S)$ and $S = B_w(s)$, defines the socially optimal amount of false positives (type I errors) and false negatives (type II errors). Relative to this benchmark, decentralized organizations clearly entail inefficiencies, given that each player does not internalize the payoff of the other, and the evaluator also neglects the research cost. When $S < S_w^*$, approval is socially excessive and the planner incurs too many false positives. Similarly when $s > s_w^*$, rejection is socially excessive and the planner incurs too many false negatives.

The starting point of our welfare analysis is the case with costless research and no discounting, corresponding to KG's setting. In this frictionless environment, Wald's planner accepts or rejects only when the uncertainty about the state of the world is completely eliminated, $(s_w^* = 0, S_w^* = 1)$, so that neither false positives nor false negatives are incurred. Given that research carries no direct cost but entails a positive option value, rejection is socially efficient under all organizational forms ($s_w^* = s^i = s^e = s^N = 0$), with false negatives never occurring. Informer authority results in socially excessive approval and too many false positives, given that the informer-authority outcome $(s^i = 0, S^i = \hat{q}_e)$ lies below the planner's bliss point. No commitment and evaluator commitment, instead, attain the first best $S^e = 1$, and thus dominate informer authority.

Once we introduce discounting, the socially efficient amount of false positives is no longer zero. Indeed, it is optimal for the planner to accept before full information is attained, so as to reduce the opportunity cost of delaying the decision. As in the frictionless environment, informer authority always results in earlier approval and more false positives than evaluator commitment. When discounting is sufficiently high, the increase in false positives associated to informer authority is socially desirable, so that informer authority yields greater social welfare than evaluator commitment, overturning the frictionless comparison.

When research is costly, another force comes into play that favors informer authority. Relative to evaluator commitment, informer authority reduces false negatives toward the socially optimal level. Intuitively, allocating authority to the informer alleviates the free-rider problem and thus incentivize

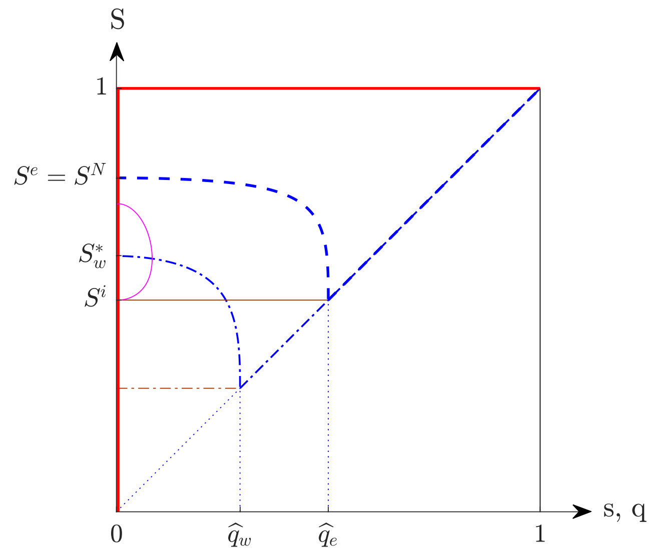

Figure 5: Welfare comparison with costless research, $c = 0$.

research, as it is socially desirable. Overall, once we introduce both discounting and research costs, informer authority dominates evaluator commitment for a large set of parameters, as we show in the following analysis.

3.1 Costless Research with Discounting

Consider first the case with costless research, $c = 0$, but positive discounting, $r > 0$, where the planner and the evaluator differ only in their approval payoffs. Because of the addition of the informer's payoff, the planner benefits from approval more than the evaluator in both states, so that the planner's myopic cutoff $\hat{q}_w$ lies to the left of $\hat{q}_e$, as shown in Figure 5.[14] More generally, the planner accepts earlier than the evaluator, $B_w(s) \leq B_e(s)$, and strictly so for any s below $\hat{q}_e$. With costless research, the planner's upper best reply, dashed and dotted in the figure, resembles the evaluator's upper best reply. It is decreasing in s, below the myopic cutoff (for $s < \hat{q}_w$) and lies along the diagonal, $B_w(s) = s$, above the myopic cutoff (for $s > \hat{q}_w$), when information has negative social value.

As r increases, the upper best replies of planner and evaluator become flatter, so that both $(s^*_w = 0, S^*_w)$ and $(s^e = 0, S^e)$ move downward along the vertical axis. In the limit, as r goes to infinity, $(s^*_w = 0, S^*_w)$ converges to $(s^*_w = 0, S^*_w = \hat{q}_w)$, while $(s^e = 0, S^e)$ approaches the informer-authority outcome $(s^i = 0, S^i = \hat{q}_e)$, which instead is independent of r. We have:

[14] Assuming that $v_i < -v^B_e$ so that information has social value, we have that $\hat{q}_w$ is still interior.

Proposition 4 *If research carries no direct cost,* $c = 0$*:*

(a) *False positives never result in all organizations,* $s_w^* = s^i = s^e = s^N = 0$*.*

(b) *No commitment and evaluator commitment coincide. Furthermore, approval is socially insufficient under both no commitment and evaluator commitment,* $S_w^* \leq S^e = S^N$*.*

(c) *There exists* $\bar{r}$ *such that:* ***(i)*** *if* $r < \bar{r}$*, approval under informer authority is socially excessive:* $S^i \leq S_w^* \leq S^e = S^N$*; and* ***(ii)*** *if* $r > \bar{r}$*, approval under informer authority is socially insufficient:* $S_w^* \leq S^i \leq S^e = S^N$*.*

(d) *There exists* $\hat{r} < \bar{r}$ *such that:* ***(i)*** *if* $r < \hat{r}$*, evaluator commitment welfare dominates informer authority; and* ***(ii)*** *if* $r > \hat{r}$*, informer authority welfare dominates evaluator commitment.*

In analogy with the frictionless environment, part (a) establishes that all decentralized institutions induce socially efficient research at the lower end. Clearly, with $c = 0$, the informer never quits researching, regardless of the organizational structure ($s_w^* = s^i = s^e = s^N = 0$), so that the planner never incurs false negatives, as it is socially optimal.

Absent a need for the evaluator to induce research at the bottom, no commitment and evaluator commitment coincide, thereby yielding the same social welfare, as shown in part (b). Since $r > 0$, both no-commitment and evaluator-commitment outcomes lie above the planner's bliss point $(s_w^* = 0, S_w^*)$, as in Figure 5. Crucially, because $S_w^* < S^e = S^N$, both no commitment and evaluator commitment result in socially insufficient approval and too few false positives, in contrast to the frictionless world where both institutions achieve the first-best.

Turning to informer authority, part (c) establishes that, with sufficient discounting, even informer authority induces socially insufficient approval and thus too few false positives. The basic reason for this distortion is that the evaluator does not take into account the informer's positive benefit from approval and thus tends to approve too little.

Finally, part (d) shows that discounting plays a key role in the welfare comparison between informer authority and evaluator commitment. When r is sufficiently low, the informer-authority outcome lies below the planner's bliss point $(s_w^* = 0, S_w^*)$, resulting in too many false positives. In this case, the welfare comparison is a priori unclear, given that the false positives are insufficient under evaluator commitment, but excessive under informer authority. As long as $r < \hat{r}$, evaluator commitment is close enough to the planner's bliss so that it still dominates informer authority, as in the absence of discounting. Crucially, as soon as r rises above $\hat{r}$, informer authority becomes welfare superior, as in Figure 5, overturning the frictionless comparison.[15] Intuitively, when the discount

[15] Graphically, the planner's iso-payoff that goes through S^i (the only iso-payoff curve we plot in the figure) is closer

factor becomes sufficiently large, the planner benefits from reducing the cost of delaying the decision by allocating authority to the informer, who is more biased toward adoption than the evaluator. As soon as r exceeds $\bar{r}$, the planner's bliss point falls below the informer-authority outcome, which in turn is below the evaluator-commitment outcome. In this case, false positives are inefficiently too few under informer authority, but even fewer under evaluator commitment, so that informer authority clearly dominates.

3.2 Costly Research

We now turn to the more general case with $c > 0$. The planner and the evaluator now differ not only in their approval payoffs, but also in that the planner takes into account the direct cost of research. As shown in Figures 6 and 7, the planner's upper best reply is now hump shaped for $s < \hat{q}_w$ and approaches 0 as s goes to 0.[16] In particular, $B_w(s)$ is increasing for $s \leq s_w^*$, reaches its maximum at the planner's stand-alone solution, s_w^*, and is then decreasing for $s \in (s_w^*, \hat{q}_w)$. When the rejection standard is excessively loose ($s < s_w^*$), an increase in s raises the marginal value of information at the top, to which the planner responds by increasing the approval standard toward S_w^*; in this case information at the bottom and at the top are strategic substitutes. When the rejection standard is excessively tough ($s > s_w^*$), the planner responds to a higher rejection standard by decreasing the approval standard away from S_w^*; in this case information at the bottom and at the top are strategic complements. Above the myopic cutoff, for $s > \hat{q}_w$, when information has negative social value, it is optimal for the planner to stop information acquisition; the upper best reply then lies along the diagonal, $B_w(s) = s$, as in the case with costless research.

Turning to the lower best replies, the different shapes of the curves are explained by one key difference between the planner and the informer: the planner's payoff, $v_w^\omega = v_e^\omega + v_i$, is state dependent because of the addition of the evaluator's approval payoff. Thus, unlike the informer, the planner values research at the bottom per se, to the extent that it provides information about the underlying state of the world. For $S < \hat{q}_w$, when the acceptance standard is below the myopic cutoff, information has negative social value so that it is optimal for the planner to forestall research; the planner's lower best reply then lies along the diagonal, $b_w(S) = S$. For $S > \hat{q}_w$, the planner's lower best reply

to the planner's bliss point (s_w^*, S_w^*) than the iso-payoff that goes through S^e. Informer authority then dominates evaluator commitment.

[16] As the rejection standard decreases, the planner expects to incur higher costs of research at the bottom. When the lower standard s is sufficiently low, the expected research cost is so high that the planner sets the approval standard $B_w(s)$ below the myopic cutoff. In the extreme case of $s = 0$, the planner finds it optimal to forestall immediately research by setting $B_w(s) = 0$.

becomes hump shaped: $b_w(S)$ is decreasing for $S \in (\hat{q}_w, S_w^*)$, increasing for $S \geq S_w^*$, and approaches 1 as S goes to 1.[17]

As illustrated in Figures 6 and 7, the informer's lower best reply (solid curve) crosses the planner's lower best reply (dashed and dotted curve) at the unique interior point $S = \hat{q}_e$, at which the informer gets the entire social payoff. To see this, note that the lower best reply for the social planner is characterized by

$$\beta_1(s,S)\, v_A^w = \beta_2(s,S)\, c/R,$$

where $v_A^w = v_i + \; S v_e^G + (1-S) v_e^B$, the expected payoff from approval of the social planner at S, replaces v_i in the analogous equation (1) for the informer. Given that $v_A^w \geq v_i$ if and only if $S \geq \hat{q}_e$, we have that the planner's lower best reply crosses the informer's lower best reply at the evaluator's myopic cutoff, $S = \hat{q}_e$, and lies to the left (respectively right) of $b_i(S)$ for $S > \hat{q}_e$ (respectively $S < \hat{q}_e$). A key consequence of this observation is that the informer-authority outcome, $(s^i = b_i(\hat{q}_e), S^i = \hat{q}_e)$, is at the intersection between the lower best replies of planner and informer.[18]

We now establish that informer authority is welfare superior to the other organizational structures in a large set of situations:

Proposition 5 ***(a)*** *False negatives are socially excessive in all decentralized organizations we consider,* $s_w^* \leq s^i \leq s^e \leq s^N$.

(b) *Evaluator commitment Pareto dominates no commitment.*

(c) *Approval is socially insufficient in all organizations,* $S_w^* \leq S^i \leq S^e \leq S^N$, *if any of* r , v_i *or* c *is sufficiently high.*

(d) *For any belief* q_0 *there exist* $\underline{q}$, $\overline{q}$, $\overline{v}_i(q_0)$ *and* $\overline{c}(q_0)$ *such that:*

(i) *if* $q_0 < \underline{q}$ *or if* $q_0 > \overline{q}$, *informer authority is welfare superior to evaluator commitment and is strictly superior for* v_i *large;*

(ii) *if* $v_i \geq \tilde{v}_i$, *informer authority is strictly welfare superior to evaluator commitment;*

(iii) *if* $c > \tilde{c}$, *informer authority is welfare superior to evaluator commitment and is strictly superior for* v_i *large.*

[17] When the approval standard is excessively lenient ($S < S_w^*$), an increase in S raises the marginal value of information at the top, thus the planner responds by reducing the rejection standard toward s_w^*; in this case information at the top is a complement with information at the bottom. When the approval standard is excessively tough ($S > S_w^*$), the planner responds to a higher approval standard by increasing the rejection standard away from s_w^*; in this case information at the top is a substitute with information at the bottom.

[18] By definition of the myopic cutoff, at $S = \hat{q}_e$ the evaluator obtains a zero payoff upon approval, so at that point the informer obtains the entire social payoff; thus $b_i(S = \hat{q}_e) = b_w(S = \hat{q}_e)$. Intuitively, if $S > \hat{q}_e$ approval benefits the evaluator and thus information carries a corresponding positive option value which the social planner should internalize; the social planner then desires to carry out more research at the bottom than the informer, $b_w(S) < b_i(S)$. Conversely, $b_w(S) > b_i(S)$ for $S < \hat{q}_e$.

According to the key result in part (a), the decentralized interaction between informer and evaluator—regardless of the specific organizational form—hinders research at the bottom and leads to excessive rejection. Even the slightest cost of research destroys Proposition 4.a's result of socially efficient false negatives under all organizational structures. To see why, recall that the evaluator has veto power over the approval decision under both informer authority and evaluator commitment. It follows that approval is granted only if the belief is above the myopic cutoff ($\hat{q}_e = S^i \leq S^e \leq S^N$), in which case a positive approval payoff accrues to the evaluator. Research thus carries a positive option value which the informer does not internalize, thereby undertaking insufficient research, $s_w^* \leq s^i \leq s^e$.[19] Because the evaluator free rides on the costly information provided by the informer, the amount of false negatives is always socially excessive.[20] In addition, part (a) crucially shows that informer authority performs better than evaluator commitment at the bottom, $s_w^* \leq s^i \leq s^e$. Intuitively, allocating authority to the informer alleviates the free-rider problem and thus fosters research at the bottom, as it is socially desirable.

Part (b) then shows that evaluator commitment Pareto dominates no commitment. The result follows directly from Proposition 3's observation that the evaluator commits to a more lenient approval standard in order to spur the informer's incentives to conduct research. Thus, relative to the no-commitment solution, not only does the evaluator benefit by revealed preference, but also the informer gains from the increased probability of acceptance. No commitment is necessarily socially inefficient and is welfare dominated by evaluator commitment, in contrast with the costless research case where the two coincide.

Part (c) focuses on the institutions' performance at the top, which depends on the specific parameters of the model. If either r, v_i or c are sufficiently high, all organizations result in too few false positives, with informer authority performing better than both evaluator commitment and no commitment. For these parameters, the preference of the planner and the informer become increasingly aligned and the veto power awarded to the evaluator particularly costly. Given the result in part (a), informer authority is closer than evaluator commitment to the planner bliss in both dimensions (s and S). In this case, informer authority is clearly welfare superior to evaluator commitment.

In general, the comparison between informer authority and evaluator commitment is more subtle and depends on the informer's motivation v_i, as well as on the cost of research and on the initial belief, as shown in part (d). We find that informer authority dominates evaluator commitment for a

[19] One additional element is also required to prove this result: the fact that for $S > \hat{q}_e$, the minimum of the lower best reply is obtained for $s = s_w^*$.

[20] As shown in Section 4, this key result can be overturned if instead the informer were to play with the planner (who internalizes the informer's decision payoff as well as the research cost).

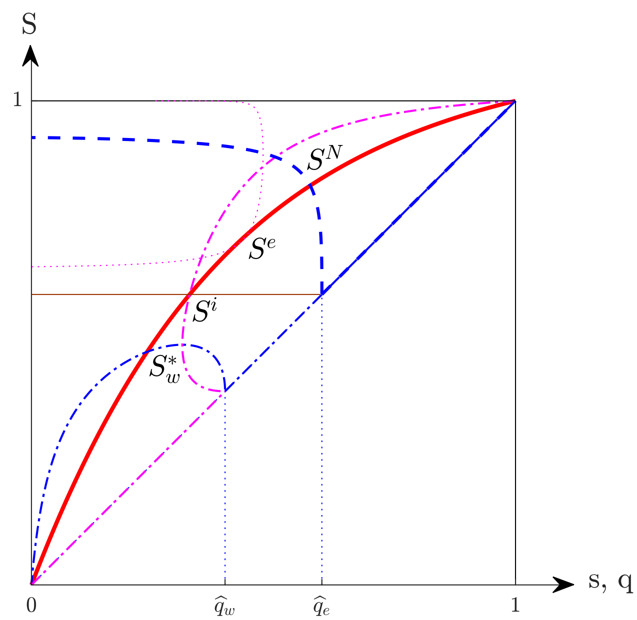

Figure 6: Highly motivated agent.

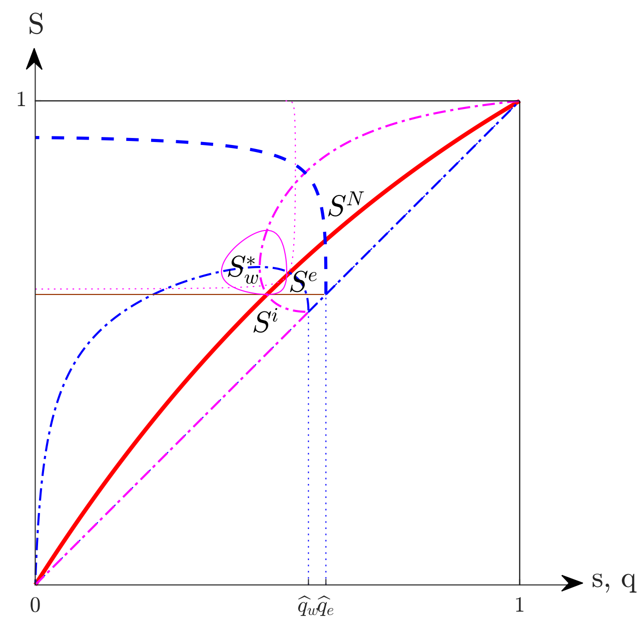

Figure 7: Poorly motivated agent.

relatively large set of parameters. To see why, note that if v_i is too low, the informer undertakes no research under both organizational structures, which then yield the same outcome and therefore the same social welfare. Only if v_i is sufficiently high the comparison between informer authority and evaluator commitment becomes meaningful. However, when the motivation v_i is large enough, the planner's incentives tend to be aligned more with the incentives of the informer than with those of the evaluator. From a social welfare standpoint, informer authority then tends to be welfare superior to evaluator commitment, as established in part (d).

In particular, when q_0 is smaller than $\underline{q}$, as soon as the motivation is large enough to incentivize research, the players find themselves in a regime represented in Figure 6, where $v_i > \overline{v}_i(q_0)$ and informer authority is socially preferred.[21] In this case, evaluator commitment never dominates informer authority, in contrast with KG's frictionless environment, where instead evaluator commitment always dominates informer authority. Figure 7 displays the boundary case where v_i is exactly equal to $\overline{v}_i(q_0)$, the threshold value defined in the proposition. As illustrated in the figure, (s^i, S^i) and (s^e, S^e) lie on the same iso-payoff curve and the two forms of commitment yields exactly the same social welfare. When the belief is intermediate ($\overline{q} < q_0 < \underline{q}$), evaluator commitment strictly dominates when v_i is below $\overline{v}_i(q_0)$ but sufficiently close to it so as not to foreclose research. In this case,

[21]Formally, the minimum requirement $\overline{v_i}(q_0)$ is the value of v_i such that $b_i(\hat{q}_e) = q_0$. We can then define $\underline{q}$ as the value of q_0 such that, for $v_i = \overline{v_i}(\underline{q})$, $S^*_w = \hat{q}_e$, that is b_i, b_w and B_w all cross at the same point. In this case, for $q_0 = \underline{q}$, if $v_i = \overline{v_i}(q_0)$, informer authority gives the first best. For values of $q_0 < \underline{q}$ (case represented in Figure 6), if $v_i > \overline{v_i}(q_0)$, informer authority strictly dominates.

informer authority dominates evaluator commitment at the bottom (where it results in less rejection, as it is socially desirable) but is dominated at the top (where it results in less approval, undesirable).

4 Approval Regulation

We now turn to our first application to approval regulation. For concreteness, consider the market for drugs where a pharmaceutical company, seeking regulatory approval to sell a drug, plays the role of informer. Over the last century, the process of drug approval in the US evolved across the different organizational forms we presented so far:[22]

- *Informer authority* characterized the early days when regulatory bodies did not exist. Patients and doctors had to decide what drugs to adopt based on the evidence voluntarily provided by the pharmaceutical companies.

- In 1905, the American Medical Association (AMA) formed the Council of Pharmacy and Chemistry. Pharmaceutical companies requesting inclusion of their drugs in a list of New and Non-Official Remedies maintained by the Council had to submit information, based on which "the Council will accept, reject or hold for further consideration", as described in the statutes of the Council.[23] This process corresponds to our *no-commitment* game.

- After the FDA was created in 1927,[24] its prerogatives were gradually increased and with the 1962 Drug Amendments the "FDA was given the authority to set standards for every stage of drug testing from laboratory to clinic",[25] creating the conditions for individualized *evaluator commitment*.

[22] See Junod (2008) for a more extensive historical account.

[23] See Cushny et al. (1905).

[24] The FDA, officially created in 1927, slowly took over from the AMA Council for Chemistry and Pharmacy and gradually increased its regulatory powers, each evolution typically reacting to a resounding scandal. The 1938 Food, Drug and Cosmetic Act required that research results be submitted to obtain approval, although the FDA had little power to mandate further research if the initial evidence was unsatisfactory. This important step made the New Drug Application (NDA) necessary to obtain approval, a procedure that still exists today. The 1962 Drug Amendments put in place a system of pre-clinical testing notification so that regulators could judge whether it was safe to start testing on humans.

[25] See Junod (2008). In its guidelines for meetings, the FDA (2009) specifies clearly that a meeting with the sponsor of the study can occur before the beginning of each phase. One of the stated goals is to define jointly the design of the future trials and the endpoints. For instance at the end of phase 2, the goal is to plan "phase 3 trials identifying major trial features such as trial population, critical exclusions, trial design (e.g., randomization, blinding, choice of control group, with explanation of the basis for any noninferiority margin if a noninferiority trial is used), choice of dose, primary and secondary trial endpoints."

For this application we need to modify our model in two ways. First, in the case of drug approval—as well as in any situation with trade—we need to account for the price $P \in \mathbb{R}$ paid to the informer by the evaluator.[26] The approval payoffs are then $v_{i,P} = v_i + P$ for the informer and $v_{e,P}^{\omega} = v_e^{\omega} - P$ for the evaluator in state ω.[27] In this section we take the price P as exogenous, given that typically government agencies that decide on approval do not set drug prices; see Scott Morton and Kyle (2011), Malani and Philipson (2012), and World Health Organization (2015).[28] Being a pure transfer among players, P does not change the payoff of the social planner; Proposition 5's results remain valid, once P is added to the threshold $\tilde{v}_i$.

Second, the identity of the player acting as evaluator has changed over time. Initially, the evaluator was the doctor or the association of doctors (AMA), whose interests could be taken to be aligned with their patients, at least to a first approximation. Over time, the role of evaluator was taken up by the FDA, a regulatory agency that should naturally aim at maximizing social welfare. In the context of our model, when does the planner—who does not have full commitment—benefit from taking up the role of evaluator? Next, we characterize the impact on social welfare of this change in the identity of the evaluator, focusing on informer authority.

4.1 Delegation of Evaluation

When taking up the role of the evaluator, the planner has an approval standard that is too lenient from the social viewpoint under informer authority. The evaluator does not internalize the informer's payoff and thus has a tougher approval standard than the planner. The planner may thus benefit from delegating play to the tougher evaluator. However, the evaluator should not be too tough either, because the increase in the upper acceptance standard in turn discourages the informer from undertaking research at the bottom, thus increasing false negatives and reducing welfare through this channel.

We explore the tradeoff between these two effects depending on the price $P \in \left(-v_i, v_e^G\right)$. As P increases from $-v_i$ to v_e^G, the planner's payoff, best replies $b_w(S)$ and $B_w(s)$, as well the solution (s_w^*, S_w^*) remain unaffected. On the contrary, as P increases, the informer has more incentives to conduct research and the lower best reply $b_{i,P}(S)$ shifts to the left, while the evaluator is inclined to

[26] In some applications we could have $P < 0$, if a fee is charged to the informer upon approval.

[27] The framework also allows for $v_i < 0$, representing production costs, as in the case of drug approval. For the informer to conduct some research it is clearly only essential that $v_{i,P} > 0$. Actually, the benefit v_i the drug company obtains from approval could also include reputation benefits from drug approval, generated for example by expected profits in other markets.

[28] For example, in the US drug prices are mostly negotiated by insurance companies, while drug approval is coordinated by the FDA. The case with endogenous price is covered in Section 5.1.

set a tougher standard, since the evaluator's benefit from approval decreases in P, and the upper best reply $B_{e,P}(s)$ shifts up. The next result characterizes the benefits of strategic delegation, extending a logic familiar from the literature at least since Fershtman and Judd (1987).[29] In what follows $s^{i_w}(P) = b_{i,P}(\hat{q}_w)$ denotes the rejection standard when the planner plays the information authority game and the exogenous price is P.

Proposition 6 *Under informer authority, for each q_0 there exist thresholds $\check{P}(q_0)$ and $\hat{P}(q_0)$, with $\check{P}(q_0) \leq \hat{P}(q_0)$, such that for any fixed price P:*

(i) *for $q_0 \in (s^{i_w}(P), \hat{q}_w)$, the planner strictly benefits from delegating the role of player to the evaluator;*

(ii) *for $q_0 \in [\hat{q}_w, S_w^*)$, the planner strictly benefits from delegating if $P \in (\check{P}(q_0), \hat{P}(q_0))$, and strictly prefers to retain the role of player if $P > \hat{P}(q_0)$; and*

(iii) *for $q_0 \in [S_w^*, 1)$, the planner strictly prefers to retain the role of player if $P > \hat{P}(q_0)$.*

For very pessimistic initial beliefs, $q_0 \in [0, s^{i_w}(P)]$, the outcome of the game under informer authority is immediate rejection regardless of whether the player who interacts with the informer is the evaluator or the planner. Thus, the planner's payoff is zero in both games. In the more interesting case where $q_0 > s^{i_w}(P)$, Proposition 6 shows that delegation serves two different purposes, depending on the initial belief. First, when the initial belief is relatively low $q_0 \in (s^{i_w}(P), \hat{q}_w)$, as in case (i), delegation to the tougher evaluator allows the planner to forestall costly but socially worthless research that would be carried out at the bottom if the planner were to retain the role of player. As long as the initial belief is in the range $(s^{i_w}(P), \hat{q}_w)$, the price level P does not determine the planner's preference for delegation. Second, when the initial belief is intermediate $q_0 \in [\hat{q}_w, S_w^*)$, as in case (ii), delegation induces more research at the top (and thus increases false negatives) and this is socially beneficial provided that P is not too high.

Case (ii) is represented in Figure 8. Recall that regardless of the price P, when $q_0 \in [\hat{q}_w, S_w^*)$ and the planner plays the game, informer authority results in immediate approval, with a positive payoff given that $q_0 \geq \hat{q}_w$. The black iso-payoff lens in Figure 8 gives all (s, S) combinations that yield the planner the same payoff as immediate approval at q_0. The outcome of the game between informer and evaluator, $\left(s^i(P), S^i(P)\right)$, lies at the intersection between the lower best replies of informer and planner, as explained in Section 3. Thus, as P increases, $\left(s^i(P), S^i(P)\right)$ moves up along the planner's lower best reply.[30] Depending on P, the outcome $\left(s^i(P), S^i(P)\right)$ then lies either inside the lens (if

[29] See also Aghion and Tirole (1996), Dessein (2002), and Armstrong and Vickers (2010) for related insights.

[30] For all S, the informer's lower best reply $b_{i,P}(S)$ increases as the price P gets higher.

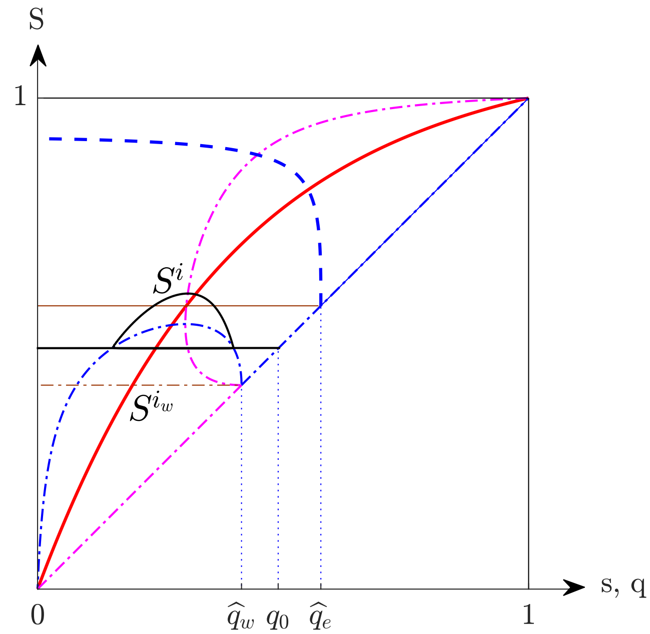

Figure 8: Value of delegation.

$P \in (\check{P}(q_0), \hat{P}(q_0)))$ or outside it ($P < \check{P}(q_0)$ or $P > \hat{P}(q_0)$). If $P < \check{P}(q_0)$, the outcome of the game under informer authority is immediate approval regardless of whether the player is the evaluator or the planner. In this case, the planner is indifferent whether or not delegating to the evaluator. If $P \in (\check{P}(q_0), \hat{P}(q_0))$, the outcome $\left(s^i(P), S^i(P)\right)$ lies within the lens and thus the planner strictly benefits from delegating to the evaluator and inducing a tougher approval standard. Finally, if $P > \hat{P}(q_0)$, the evaluator's approval standard becomes too tough; excessive research is carried out at the top and the planner then strictly prefers to retain the role of player.

As the initial belief increases toward S_w^*, the planner loses interest in conducting research. The lens becomes smaller and the interval $(\check{P}(q_0), \hat{P}(q_0))$ shrinks. When the belief becomes too optimistic, specifically when $q_0 \in [S_w^*, 1)$, as in case (iii), the lens collapses into a single point and the interval $(\check{P}(q_0), \hat{P}(q_0))$ becomes empty. Any outcome with research yields the planner a lower payoff than immediate approval. For $q_0 \in [S_w^*, 1)$, we then have a special instance of case (ii) for which $\check{P}(q_0) = \hat{P}(q_0)$, so that the planner never strictly benefits from delegation.

4.2 Planner Commitment

As discussed in the previous section, delegation to the evaluator may be welfare improving when authority is allocated to the informer. By revealed preferences, though, delegation is unambiguously suboptimal when the player in charge of the approval decision has commitment power. Based on this observation and building on the simpler analysis of evaluator commitment in Section 2.3, we now characterize the planner's commitment solution $S^{w_w}(q_0)$ for the game between the planner and the informer.

In Section 2.3, we showed that the evaluator optimally commits to an approval standard below the Nash level, in order to encourage the informer to undertake research. The evaluator always benefits from the additional research at the bottom induced by the lenient commitment, given that the cost of research is borne only by the informer. The planner, instead, also cares about the research cost. Hence, when the informer is particularly motivated to carry out research, for instance because of a high price, it is optimal for the planner to commit to a level above Nash so as to discourage the informer from researching.

The planner's optimal commitment $S^{w_w}(q_0)$ is characterized as follows:

Proposition 7 *If the price paid to the informer is relatively low, $P < \bar{P}$, there is a threshold of the initial belief $\tilde{q} \in (s_w^*, \hat{q}_w)$, which depends on P, such that:*

(a) *For $q_0 \in (0, \tilde{q}]$, the planner blocks research by choosing a sufficiently high standard: $S^{w_w} > b_i^{-1}(q_0)$.*

(b) *For $q_0 \in (\tilde{q}, S^{N_w})$, the planner chooses an interior commitment such that:*

(i) *if $P \in (0, P^*)$, then $S^{w_w}(q_0)$ is increasing in q_0 and below the Nash level, $S^{w_w}(q_0) < S^{N_w}$;*

(ii) *if $P \in (P^*, \underline{P})$, then $S^{w_w}(q_0)$ is decreasing in q_0 and above the Nash level, $S^{w_w}(q_0) > S^{N_w}$, with no discontinuity at $\tilde{q}$: $S^{w_w}(\tilde{q}) = b_i^{-1}(\tilde{q})$; and*

(iii) *if $P \in (\underline{P}, \bar{P})$, then $S^{w_w}(q_0)$ is decreasing in q_0 and above the Nash level, $S^{w_w}(q_0) > S^{N_w}$ with a downward discontinuity at $\tilde{q}$: $\lim_{q_0 \to \tilde{q}^+} S^{w_w}(q_0) < b_i^{-1}(\tilde{q})$.*

(c) *For $q_0 > S^{N_w}$, the planner approves immediately: $S^{w_w}(q_0) \leq q_0$.*

For low initial beliefs, $q_0 < \tilde{q}$, as in case (a), any commitment that induces research yields a negative social payoff. Thus, the planner optimally commits to a sufficiently high approval standard, so as to forestall costly information acquisition. Intuitively, when the initial belief is sufficiently unfavorable, the expected cost of research exceeds the corresponding social benefits, so that the planner benefits from a blocking commitment that induces the informer to abandon research. In some circumstances, the planner blocks research that would be carried out under the other organizational forms. Importantly, this cannot be the case under evaluator commitment, given that the evaluator never gains from curbing the informer's incentive to conduct research.

When the optimal commitment is interior, as in case (b), the shape of the commitment path crucially depends on the value of P. For low prices, as in case (i), the informer undertakes insufficient research at the Nash outcome relative to the social optimum. Therefore, the planner obtains a first-order gain by committing to an approval standard below Nash so as to encourage research, similar to the evaluator-commitment outcome. We have $\lim_{q_0 \to S^{N_w}+} S^{w_w}(q_0) = S^{N_w}$, with the planner commitment solution converging to the Nash level along an upward sloping path, as shown in Figure 10 in Appendix A.

When the price is raised to the critical level P^*, the Nash solution is socially optimal, so that commitment has no value for the planner. For any price above this level, as in cases (ii) and (iii), the Nash outcome results in excessive research at the bottom (and thus in insufficient false negatives) relative to the social optimum. So as to discourage research, the planner thus optimally commits to an approval standard above Nash. As $q_0 \to S^{N_w}$, the commitment solution now moves along a downward sloping path that still converges to the Nash outcome. In case (ii) this path is continuous, as shown in Figure 11 in Appendix A.

Case (iii) highlights a discontinuity in the commitment path when $P \in (\underline{P}, \bar{P})$. For P slightly above $\underline{P}$, the informer's lower best reply (red curve in Figure 9) crosses the planner's zero-level iso-payoff at $q_0 = \check{q}$ (dashed black), where $\check{q} := b_i(\check{S}) = b_w(\check{S})$ is defined as the point of intersection between the lower best replies of the two players. At this level of the initial belief, thus, the planner chooses an interior commitment that yields a strictly positive payoff. Therefore, there must be a belief $q_0 = \tilde{q} < \check{q}$ for which the planner obtains exactly zero at the optimal interior commitment. The corresponding iso-payoff (continuous black) is tangent to the informer's lower best reply at a belief, $b_i(S^{w_w}(\tilde{q}))$, which is strictly lower than $\tilde{q}$, as illustrated in Figure 9. At $q_0 = \tilde{q}$, the planner is therefore indifferent between choosing the interior commitment, $S = S^{w_w}(\tilde{q})$ and a blocking commitment, corresponding to any $S \geq b_i^{-1}(\tilde{q})$. Since $b_i(S^{w_w}(\tilde{q})) < \tilde{q}$, such an interior commitment induces the informer to set the rejection standard below $\tilde{q}$. The planner's optimal interior commitment $S^{w_w}(\tilde{q})$ must be strictly below the lowest blocking commitment $b_i^{-1}(\tilde{q})$ and this gives rise to the discontinuity displayed in the figure.[31]

These results shed light on approval regulation in the pharmaceutical sector. As outlined above, social welfare surely increased when decision-making authority was transferred from the AMA (representing doctors and thus indirectly the interest of patients) to the FDA (a government agency

[31] To understand the discontinuity, consider the dashed curve in Figure 9, which represents the value function of the planner, for $S = S^{w_w}(\tilde{q})$, as a function of the initial belief q_0. As illustrated in the figure, at $q_0 \leq b_i(S^{w_w}(\tilde{q}))$ and $q_0 = \tilde{q}$, the value of the planner equals zero. However, when the initial belief is between these two points, the optimal commitment at $\tilde{q}$, yields a strictly negative value. Indeed, for $q_0 \in (b_i(S^{w_w}(\tilde{q})), \tilde{q})$ the dashed curve lies below the horizontal axis so that for those levels of the initial belief the planner would prefer to block research rather than adopting an interior commitment. Since $b_i(S^{w_w}(\tilde{q})) < \tilde{q}$, such range of beliefs is not empty so that a discontinuity arises.

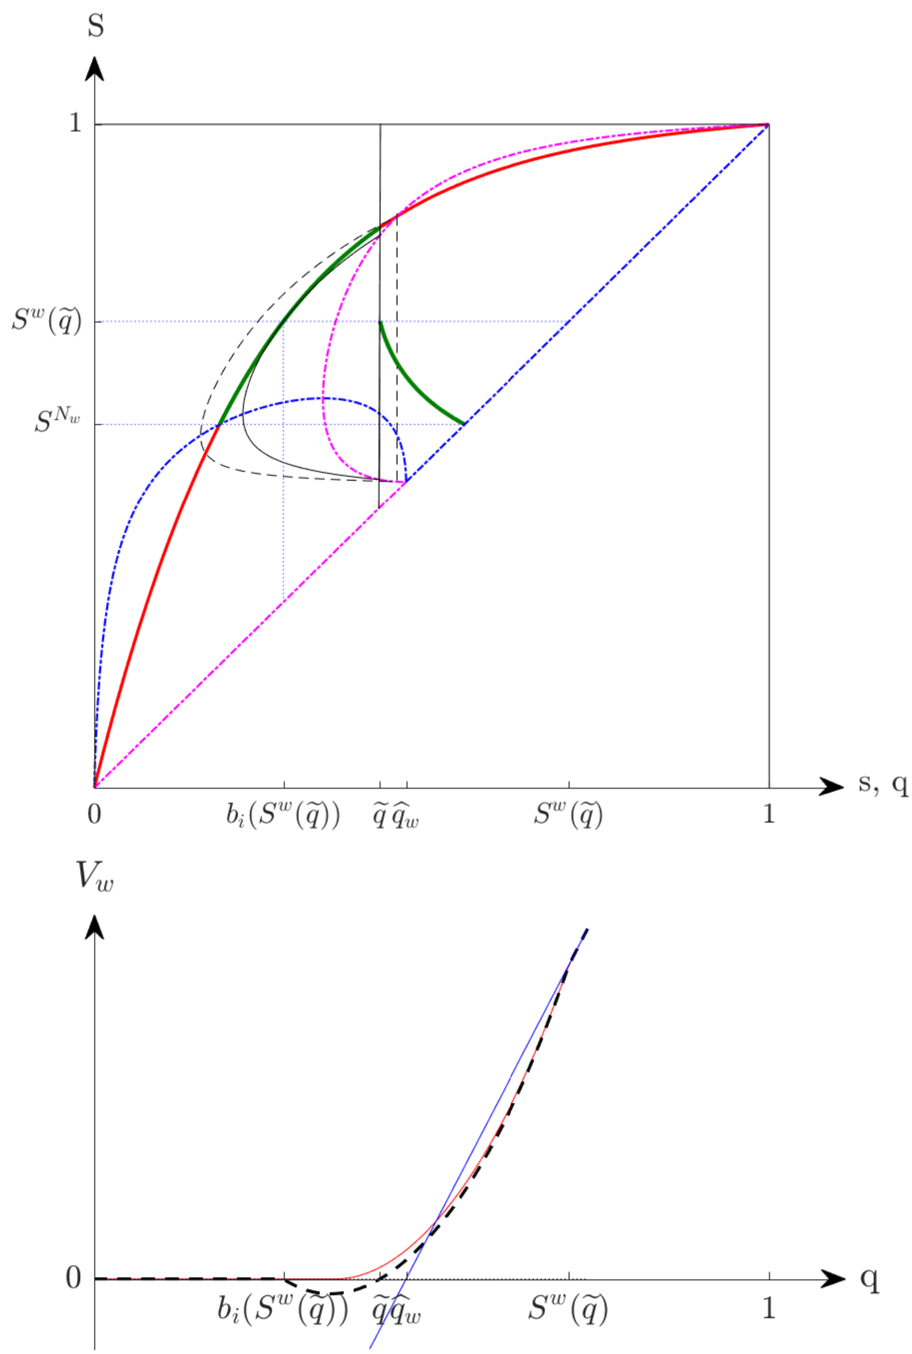

Figure 9: (a) Discontinuous commitment path; (b) Planner's value function.

that should act in the collective interest) that has the power to enact planner commitment. In some circumstances, the FDA chooses a blocking commitment, along the lines of case (a). Before being allowed to start clinical trials on humans, pharmaceutical firms need to submit an Investigational New Drug (IND) application that must include results of tests on animals or clinical tests on humans performed abroad. Based on this initial evidence, which can be interpreted as q_0 in our model, the FDA decides to allow or put on hold the clinical testing.[32]

If clinical testing is allowed, requirements partly depend on the market potential of the drug. For most drugs, pharmaceutical companies can extract high prices upon approval. Tough commitment by the FDA, as in cases (ii) and (iii), thus serves the purpose of discouraging pharmaceutical companies from running excessive clinical trials. These trials are socially costly and pharmaceutical companies have excessive incentives to undertake research to prove the drug's effectiveness. However, there are certain categories of drugs where the planner commits to a lower standard of approval. In the case of the orphan drugs, for instance, the price the pharmaceutical companies can extract is small and the planner is mostly concerned about providing incentives to conduct research, as in case (i).[33]

5 Technology Transfer

So far we considered environments where the transfer price P upon approval was exogenously set. We now turn to unregulated market settings. We begin by considering a setting with an informer who sells to a single buyer/evaluator, to then turn to a setting with competing evaluators.

5.1 Endogenous Price

Recall from Section 4 that the price P affects the incentives of both informer and evaluator, but leaves the social planner's bliss point (s_w^*, S_w^*) unchanged. Note that the social planner could always achieve the first-best Wald outcome by a judicious choice of the price P. To fix ideas, consider the informer-authority solution in our baseline game between informer and evaluator. There always exists a transfer $P^*(v_i)$ such that the first best (s_w^*, S_w^*) is achieved under informer authority. Furthermore, there exists a threshold $\hat{v}_i$ such that $P^*(v_i) \geq 0$ (the optimal transfer is from evaluator to informer) if and only if $v_i \leq \hat{v}_i$ (informer is poorly motivated).

[32] Lapteva and Pariser (2016) report that out of the 1410 applications received in 2013, 8.9 per cent were put on hold and half of them were eventually authorized.

[33] Orphan drugs are those that treat rare diseases considered to have a small commercial market. They often benefit from accelerated regulatory approval programs.

An increase in P increases the informer's incentives to do research and thus shifts the informer's lower best reply to the left. The informer-authority outcome S^i is the unique value of S such that $b_i(S)$ and $b_w(S)$ intersect, as explained in Section 3.2. As P increases, S^i moves upward along the social planner's lower best reply $b_w(S)$. Thus, to achieve the first best, P needs to be chosen such that S^i moves all the way to S_w^*. If the informer is highly motivated (with $v_i > \hat{v}_i$), as in Figure 6, the informer's lower best reply needs to move downward to intersect at $S = S_w^*$, thus the informer should be charged upon approval. On the contrary, a poorly motivated informer (with $v_i < \hat{v}_i$) should be paid, as in Figure 7.

While the price is naturally fixed exogenously in many applications (e.g., drug approval or publication process), in instances involving trade the transfer can be endogenously chosen by the players.[34]

Proposition 8 ***(a)*** *If the informer endogenously sets the price when approval is requested, then the-first best outcome* (s_w^*, S_w^*) *is achieved, regardless of the organizational form.*

(b) *If the evaluator endogenously sets the price when approval is requested, then no information is collected, regardless of the organizational form.*

Part (a) follows the logic of the Coase theorem. If the informer endogenously sets the price when requesting approval, the informer is able to extract the full surplus of the evaluator regardless of the organizational form. Thus, the informer's incentives become perfectly aligned with those of the social planner and the organization of persuasion is irrelevant.[35] According to part (b), the evaluator who endogenously chooses the price expropriates the informer at the time when approval is requested. Because of this hold up, the informer's incentives to collect information are eliminated altogether. This holds for all organizational structures, as long as the evaluator cannot commit on the price.

Even when the informer sets the price, trade often involves several competing evaluators trying to contract with the informer and often exerting externalities on each other when they conclude a trade. Proposition 8.a's irrelevance result then no longer holds, as we explore next.

[34] Branco, Sun, and Villas-Boas (2010), instead, characterize instead the optimal price set by a seller *before* the buyer decides to acquire information.

[35] It would be interesting to extend the analysis to the case with pre-existing private information; a natural starting point is given by Daley and Green (2012).

5.2 Competing Evaluators

Focus on a developer selling a technology to competing producers and running tests on the technology to convince the potential buyers. As documented by Arora, Fosfuri, and Gambardella (2001), markets for technology in high-tech industries are often divided between upstream technology manufacturers and downstream producers. This is the case for instance in the pharmaceutical industry, where smaller biotech firms do a large share of the research and transfer their discoveries to competing large pharmaceutical firms at some point during the clinical trials process, as well as in the chemical industry with the development of specialized engineering firms; see also Gans, Hsu, and Stern (2008).[36] When one producer obtains the technology which turns out to be good, competing evaluators suffer a negative externality, due to lost profits from facing a stronger competitor.

Specifically, payoffs are as follows. First, to simplify the exposition, we assume the technology developer cares only about the price that can be obtained for the technology so that $v_i = 0$. If acquiring the technology at price P, producer k obtains payoff $v_k^G - P$ in the good state and $v_k^B - P$ in the bad state. To our baseline model we add two elements that reflect relevant features of the application. First, while none of the producers have acquired the technology, they obtain a benefit A_k from the existing technologies; we assume $v_k^B \leq A_k$ so that even if the new technology turns out not to be successful, the producer can keep using the old one. Second, if the other producer l acquires the technology and the state is good, producer k obtains benefit $D_k < A_k$ (the sale is assumed to be exclusive).[37] Thus, the fact that the other producer gets the technology imposes a negative externality on the first producer. Section 5.3 provides a particular example of this class of payoffs. We consider informer authority: if the informer proposes the technology to one producer and it is refused, the informer cannot propose it again to the same producer. We also assume that the informer endogenously chooses the price.

In this environment, the key mechanism we identify is that, if the first producer refuses, the informer can credibly threaten to conduct additional research to convince the second producer.

Proposition 9 ***(a)*** *The lower and upper best replies of the informer decrease with the externality.*

(b) *Both producers are worse off than if the new technology did not exist.*

(c) *When the two producers differ only in the value of the technology in the good state, $v_1^G > v_2^G$, if D_j is small enough, the developer sells the technology to the less efficient producer 2.*

[36]Similarly, managers within a firm compete when examining ideas of their employees. And within academia, if a scientific journal rejects a paper, competing journals in the field will have the chance to publish an improved version of the article.

[37]Many contracts in the market for technology are exclusive, as in the case of transfer of patent ownership.

According to result (b), both evaluators are made worse off by the introduction of the new technology. The sequential nature of research is essential for this result. When facing the first producer, the informer can threaten, in case of rejection, to go and see the other producer. Before approaching the second producer, the informer does more research. Following rejection, the first producer thus runs the risk of suffering a negative externality, in case the new information is sufficiently favorable to make the second producer accept the project and the project turns out to be good. Anticipating this, the first producer is thus ready to accept a project even though acceptance results in a payoff reduction relative to the absence of the new technology.

Our model allows us to characterize the extent of the loss for the producers. Consider the subgame following a refusal by producer k, where the informer needs to determine for what belief and at what price to offer the project to the other producer. The informer in this subgame, consistent with the Coasian result of Proposition 8, extracts the full surplus of producer l so that the informer acts as a social planner. The outcome is denoted (s_l^*, S_l^*), corresponding to producer l's bliss point.

Consider now the offer made to the first producer. If the informer decides to make an offer to producer k at belief q, the informer charges the maximum price that the producer is ready to accept, denoted $P(q)$, which is given by the indifference condition

$$qv_k^G + (1-q)v_k^B - P(q) \;=\; A_k \;-\; q\,\phi(q, s_l^*, S_l^*)(A_k - D_k),$$

where ϕ, defined in equation (11) in Appendix A, measures the expected discounted probability that the informer obtains evaluator l's approval by reaching S_l^* before s_l^* starting from q, conditional on the state being G. The choice of when to stop research will be partly determined by the size of the externality the informer can extract at that belief.

Part (a) shows that the presence of the externality affects the best replies. As the size of the externality increases (decrease in D_k), the payoff the informer expects in the good state increases, so the informer is willing to do more research at the bottom (the lower best reply decreases). For the upper best reply, two countervailing effects are at play. On the one hand, the same increase in the payoff in the good state makes the informer more eager to accept early and not delay this payoff. On the other hand, delaying acceptance can lead to the accumulation of more positive evidence that can allow the informer to extract a higher externality (i.e., $\phi(q, s_l^*, S_l^*)$ is increasing in q). Proposition 9.a shows that the first effect dominates so that the upper best reply decreases as the size of the externality increases.

Part (c) considers the case in which the two producers differ only in the value they derive from the new technology in the good state and characterize when the informer in equilibrium ends up selling the technology to the *less* efficient producer. The idea is the following. On the one hand, the first producer (more efficient) obtains a higher immediate payoff from accepting the project since v_1^G is higher. On the other hand, the first producer faces a smaller pressure from the competitive externality, since following refusal, the informer needs to carry out additional research to convince the second producer. Indeed, producer 2's bliss point is such that $s_1^* < s_2^*$ and $S_1^* < S_2^*$, so that $\phi(q, s_1^*, S_1^*) > \phi(q, s_2^*, S_2^*)$ for all values of q. Since these expected discounted probabilities are independent of the size of the externality $A - D$, if the externality is large enough, the second effect dominates and the informer prefers selling the technology to the second producer who suffers from a higher pressure due to the externality effect.

5.3 Revisiting the Persistence of Monopoly

To conclude, we consider in this section the application of our results to the special case where the two evaluators are an incumbent monopolist and a potential entrant, as in Gilbert and Newbery (1982) or Reinganum (1983). The monopolist is currently enjoying monopoly profits with the old technology Π_{m0}. A new technology has been developed. If the new technology is bad it generates zero profits. If it is good, it generates profits $\Pi_{m1} > \Pi_{m0}$ if the incumbent acquires it. If the entrant acquires a good technology, the entrant obtains profits Π_E but the incumbent's profits drop to Π_I, with $\Pi_{m1} \geq \Pi_I + \Pi_E$.[38] Note that for a drastic innovation, $\Pi_I = 0$ and $\Pi_{m1} = \Pi_E$.

We now study how the introduction of a technology producer affects the game between the incumbent and the entrant. To clarify the tradeoffs and facilitate the comparison with the literature, we assume the transfer is fixed at a given value P. The entrant does not suffer any externality, regardless of being approached first or second, and so accepts the project for beliefs above the myopic cutoff $\hat{q}_E$ defined by

$$\hat{q}_E = \frac{P}{\Pi_E}.$$

The incumbent, on the other hand, when approached first, accepts the project for beliefs above the competitive cutoff $\tilde{q}_I$ defined below, where this competitive acceptance cutoff depends on how likely the entrant is to end up accepting the project following rejection[39]

$$\tilde{q}_I = \frac{P}{\Pi_{m1} - \Pi_{m0} + \phi(\tilde{q}_I, b_i(\hat{q}_E), \hat{q}_E)(\Pi_{m0} - \Pi_I)}.$$

[38]This illustration maps into the above payoffs as follows: for the incumbent $v_I^G = \Pi_{m1}$, $v_I^B = \Pi_{m0}$ (retaining the old technology), $A_I = \Pi_{m0}$, and $D_I = \Pi_I$; and for the entrant $v_E^G = \Pi_E$ and $v_E^B = A_E = D_E = 0$.

[39]If the incumbent rejects, the informer performs research in the interval $(b_i(\hat{q}_E), \hat{q}_E)$.

We see that if $\phi = 1$, the incumbent is the first to be approached and acquires the technology since $\Pi_E \leq \Pi_{m0} - \Pi_I$. This corresponds to Gilbert and Newbery's (1982) result that states that the incumbent has more incentives to outbid the rival in order to preserve monopoly power. Here it requires that there be no uncertainty on the research process, so that the informer, following refusal by the monopolist, is very quickly able to reach a level of knowledge sufficient for the entrant to accept.

At the other extreme, if ϕ approaches zero, the entrant will be the first one approached and acquire the technology as long as $\Pi_E > \Pi_{m1} - \Pi_{m0}$. This corresponds to the replacement effect highlighted by Reinganum (1983): while the technology is acquired by neither of the two producers, the incumbent enjoys monopoly profits Π_{m0} and therefore has little incentive to replace this flow of profits by Π_{m1}. If the research technology of the informer is very imprecise, it will take a long time to convince the entrant to buy.

Proposition 10 *If* $\Pi_E < \Pi_{m1} - \Pi_{m0}$*, the informer approaches the incumbent first. If* $\Pi_E > \Pi_{m1} - \Pi_{m0}$*:*

(i) *When the innovation is drastic, the informer approaches the entrant first.*

(ii) *When the innovation is not drastic, there exists* $\mu^*(c)$*, such that the informer approaches the incumbent first if and only if the research process is sufficiently precise,* $\mu > \mu^*$*.*

In the setting we consider, the order of play is chosen by the technology producer, who determines whether to first approach the incumbent or the entrant. In that sense our model encompasses both Gilbert and Newbery's (1982) preemptive bidding effect and Reinganum's (1983) replacement effect. We see that a key determinant of this choice is the efficiency of the research technology captured by the parameter μ.

6 Conclusion

We conclude by discussing some of the features of our model and possible extensions. Our analysis is made particularly tractable by the assumption that the state is binary.[40] The simplifying assumption that the informer's benefit upon approval is state independent is natural in the context of the application to drug approval. FDA approval largely shelters pharmaceutical companies from liability, unless they are shown to have been hiding or misrepresenting information. Our results hold more

[40] Continuous-time models with continuous state are typically characterized numerically. However, Fudenberg, Strack, and Strzalecki (2015) recently made strides in the analytical characterization of Wald's decision-theoretic model with continuous state, opening the way for future work on strategic extensions.

generally when the informer's payoff from approval is state dependent, provided that the informer is biased toward approval relative to the evaluator.[41]

Our model can be extended to address a number of recent concerns in the context of drug approval. First, our continuous-time model can be extended to allow for withholding of negative results and costly misrepresentation, an issue regulators are currently grappling with.[42] Second, there is a tractable way to incorporate flow of information following market introduction and reversal of approval through recall on the basis of post-approval information.[43] Future work could also add competition among researchers.[44]

[41]An earlier working paper version of this work also extends the analysis to allow for rejection's payoff to be state dependent. It is also natural to allow the informer to have pre-existing private information; for a starting point see Taylor and Yildirim's (2011) static analysis.

[42]This paper's specification constrains reporting to be truthful at the moment of application, for example because misrepresentation is infinitely costly as in the disclosure models of Grossman (1981) and Milgrom (1981). Our earlier working paper characterizes the impact of ex post lying costs à la Kartik, Ottaviani, and Squintani (2007) on ex ante incentives for dynamic information collection. See also Shavell (1994), Dahm, Gonzalez, and Porteiro (2009), Henry (2009), and Felgenhauer and Schulte (2014) for work on partial disclosure of research results.

[43]For example, see Zuckerman, Brown, and Nissen (2011) on the prevalence of recalls of drugs and medical devices depending on the approval process. For economic analyses of mandatory and voluntary product recalls see Marino (1997) and Spier (2011) respectively, as well as Rupp and Taylor (2002) on their relative empirical prevalence in the car industry.

[44]In this vein, Bobtcheff, Bolte, and Mariotti (2017) focus on researchers' incentives to improve the quality of their ideas under the threat of being scooped by competing researchers—but abstracting from the quality of the evaluation process on which we focus.

References

Aghion, Philippe, and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy*, 105(1): 1–29.

Armstrong, Mark, and John Vickers. 2010. "A Model of Delegated Project Choice." *Econometrica*, 78(1): 213–244.

Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella. 2001. *Markets for Technology: The Economics of Innovation and Corporate Strategy*. MIT Press.

Bizzotto, Jacopo, Rudiger Jesper, and Adrien Vigier. 2016. "How to Persuade a Long-Run Decision Maker." Unpublished working paper.

Bobtcheff, Catherine, Jérôme Bolte, and Thomas Mariotti. 2017. "Researcher's Dilemma." *Review of Economic Studies*, forthcoming.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica*, 67(2): 349–374.

Branco, Fernando, Monic Sun, and J. Miguel Villas-Boas. 2010. "Optimal Search for Product Information." *Management Science*, 58(11): 2037–2056.

Brocas, Isabelle, and Juan D. Carrillo. 2007. "Influence through Ignorance." *RAND Journal of Economics*, 38(4): 931–947.

Brocas, Isabelle, Juan D. Carrillo, and Thomas R. Palfrey. 2012. "Information Gatekeepers: Theory and Experimental Evidence." *Economic Theory*, 51(3): 649–676.

Chan, Jimmy, Alessandro Lizzeri, Wing Suen, and Leeat Yariv. 2015. "Deliberating Collective Decisions." Unpublished working paper.

Cushny, Arthur R., C. Lewis Diehl, C. S. N. Hallberg, Robert A. Hatcher, L. F. Kebler, J. H. Long, F. G. Novy, W. A. Puckner, Samuel P. Sadtler, J. O. Schlotterbeck, Geo. H. Simmons, Torald Sollmann, Julius Stieglitz, M. I. Wilbert, and H. W. Wiley. 1905. "American Medical Association. Council on Pharmacy and Chemistry," *California State Journal of Medicine*, 3(4): 103–104.

Dahm, Matthias, Paula Gonzalez, and Nicolas Porteiro. 2009. "Trials, Tricks and Transparency: How Disclosure Rules Affect Clinical Knowledge." *Journal of Health Economics*, 28(6): 1141–1153.

Daley, Brendan, and Brett Green. 2012. "Waiting for News in the Market for Lemons." *Econometrica*, 80(4): 1433–1504.

Dessein, Wouter. 2002. "Authority and Communication in Organizations." *Review of Economic Studies*, 69(4): 811–838.

Dvoretsky, Aryeh, Jack Keifer, and Jacob Wolfowitz. 1953. "Sequential Decision Problems for Processes with Continuous Time Parameter: Testing Hypothesis." *Annals of Mathematical Statistics*, 24: 254–264.

FDA. 2009. "Guidance for Industry: Formal Meetings Between the FDA and Sponsors or Applicants." May 2009 Procedural.

Felgenhauer, Mike, and Elisabeth Schulte. 2014. "Strategic Private Experimentation." *American Economic Journal: Microeconomics*, 6(4): 74–105.

Fershtman, Chaim, and Kenneth L. Judd. 1987. "Equilibrium Incentives in Oligopoly." *American Economic Review*, 77(5): 927–940.

Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki. 2015. "Stochastic Choice and Optimal Sequential Sampling." SSRN Working Paper.

Gans, Joshua, David Hsu, and Scott Stern. 2008. "The Impact of Uncertain Intellectual Property Rights on the Market for Ideas: Evidence for Patent Grant Delays." *Management Science*, 54(5): 982–997.

Gilbert, Richard J., and David M. G. Newbery. 1982. "Preemptive Patenting and the Persistence of Monopoly." *American Economic Review*, 72(3): 514–526.

Green, Brett, and Curtis R. Taylor. 2016. "Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects." *American Economic Review*, 106(12): 3660–3699.

Grenadier, Steven R., Andrey Malenko, and Nadya Malenko. 2016. "Timing Decisions in Organizations: Communication and Authority in a Dynamic Environment." *American Economic Review*, 106(9): 2552–2581.

Grossman, Sandford J. 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics*, 24(3): 461–483.

Gul, Faruk, and Wolfgang Pesendorfer. 2012. "The War of Information." *Review of Economic Studies*, 79(2): 707–734.

Guo, Yingni. 2016. "Dynamic Delegation of Experimentation." *American Economic Review*, 106(8): 1969–2008.

Henry, Emeric. 2009. "Strategic Disclosure of Research Results: The Cost of Proving Your Honesty." *Economic Journal*, 119(539): 1036–1064.

Junod, Suzanne White. 2008. "FDA and Clinical Drug Trials: A Short History." in: *A Quick Guide to Clinical Trials*, eds. Madhu Davies and Faiz Kerimani. Washington: Bioplan, Inc.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion," *American Economic Review*, 101(6): 2590–2615.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani. 2007. "Credulity, Lies, and Costly Talk." *Journal of Economics Theory*, 134(1): 93–116.

Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73(1): 39–68.

Lapteva, Larissa, and Anne R. Pariser. 2016. "Investigational New Drug Applications: A 1-Year Pilot Study on Rates and Reasons for Clinical Hold." *Journal of Investigative Medicine*, 64(2): 376–382.

Malani, Anup, and Tomas J. Philipson. 2012. "The Regulation of Medical Products." *Oxford Handbook of the Economics of the Biopharmaceutical Industry*.

Marino, Anthony M. 1997. “A Model of Product Recalls with Asymmetric Information.” *Journal of Regulatory Economics*, 12(3): 245–265.

Milgrom, Paul R. 1981. “Good News and Bad News: Representation Theorems and Applications.” *Bell Journal of Economics*, 12(2): 380–391.

Mikhalevich, Vladimir S. 1958. “Bayesian Choice Between Two Hypotheses on the Mean Value of a Normal Process.” *Visnyk Kyivskogo Universytetu*, 1(1): 101–104.

Moscarini, Giuseppe, and Lones Smith. 2001. “The Optimal Level of Experimentation.” *Econometrica*, 69(6): 1629–1644.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov. 2016. “Persuading the Regulator to Wait.” Unpublished working paper.

Reinganum, Jennifer F. 1983. “Uncertain Innovation and the Persistence of Monopoly.” *American Economic Review*, 73(4): 741–748.

Rupp, Nicholas G., and Curtis R. Taylor. 2002. “Who Initiates Recalls and Who Cares? Evidence from the Automobile Industry.” *Journal of Industrial Economics*, 50(2): 123–149.

Scott Morton, Fiona, and Margaret Kyle. 2011. “Markets for Pharmaceutical Products.” *Handbook of Health Economics*, Volume 2: 763–823.

Shavell, Steven. 1994. “Acquisition and Disclosure of Information Prior to Sale.” *RAND Journal of Economics*, 25(1): 20–36.

Shiryaev, Albert N. 1967. “Two Problems of Sequential Analysis.” *Cybernetics*, 3(2): 63–69.

Shiryaev, Albert N. 1978. *Optimal Stopping Rules*. Springer, Berlin.

Spier, Kathryn E. 2011. “Product Safety, Buybacks, and the Post-Sale Duty to Warn.” *Journal of Law, Economics, and Organization*, 27(3): 515–539.

Stokey, Nancy. 2009. *The Economics of Inaction*. Princeton University Press, Princeton.

Strulovici, Bruno. 2010. “Learning while Voting: Determinants of Strategic Experimentation.” *Econometrica*, 78(3): 933–971.

Taylor, Curtis, and Huseyin Yildirim. 2011. “Subjective Performance and the Value of Blind Evaluation.” *Review of Economic Studies*, 78(2): 762–794.

Wald, Abraham. 1947. *Sequential Analysis*. John Wiley and Sons, New York.

World Health Organization. 2015. *2015 Global Survey on Health Technology Assessment by National Authorities*.

Zuckerman, Diana M., Paul Brown, and Steven E. Nissen. 2011. “Medical Device Recalls and the FDA Approval Process.” *Archives of Internal Medicine*, 171(11): 1006–1011.

A Appendix A: Derivations and Proofs

Log-odds Parametrization

To facilitate derivations, some of the results in the appendix are presented using the following log-odds parametrization of beliefs

$$\sigma = \log \frac{q}{1-q} \in (-\infty, \infty).$$

In this log-odds space, the lower and upper standards of research are $\left(\mathrm{s}=\ln\frac{s}{1-s}, \mathrm{S}=\ln\frac{S}{1-S}\right)$ the logit function of the standards in the regular belief space, (s,S). Finally, $\hat{\sigma}_e = \ln\frac{\hat{q}_e}{1-\hat{q}_e}$ denotes the myopic cutoff in the log-odds space.

Updating of Beliefs

If research is undertaken until time $t > 0$, the realization of the stochastic process v_t is a sufficient statistic for all the information collected until this instant of time and will be used to update beliefs. The log-likelihood ratio of observing $v_t = \gamma$ under the two states is

$$\log \frac{h\left(\frac{\gamma-\mu t}{\rho\sqrt{t}}\right)}{h\left(\frac{\gamma+\mu t}{\rho\sqrt{t}}\right)} = \frac{2\mu\gamma}{\rho^2},$$

where h is the density of a standard normal distribution. According to Bayes' rule, the log posterior probability ratio is equal to the sum of the log prior probability ratio and the log-likelihood ratio. Thus, the posterior belief at time t is $\sigma_t = \sigma_0 + \Sigma_t'$, where $d\Sigma'$ is a Wiener process with drift

$$\mu' = \frac{2\mu^2}{\rho^2} \tag{4}$$

if the state is G and $-\mu'$ if the state is B and instantaneous variance $2\mu'$. Normalizing WLOG $\rho = 1$, μ parametrizes the speed of learning.

Expected Utility in Research Region

If the upper and lower standards (s,S) are given, for $\sigma \in (\mathrm{s},\mathrm{S})$ we have that the expected utility of player j (with cost of research c_j and benefits v_j^G and v_j^B from approval) follows

$$u_j(\sigma) = e^{-rdt} E[u_j(\sigma + d\Sigma')] - c_j dt.$$

Following Stokey (2009, Chapter 5), starting in the intermediate region, we let T be the first time the belief hits either s or S. The direct monetary cost of searching is given by $\int_0^T c_j e^{-rt} dt = \frac{c_j}{r} - \frac{c_j}{r} e^{-rT}$. Once we define the expected discounted probabilities

$$\begin{aligned} \Psi(\sigma,\omega) &= E[e^{-rT} | \sigma(T) = \mathsf{S}, \omega] \ \Pr[\sigma(T) = \mathsf{S} | \omega] \\ \psi(\sigma,\omega) &= E[e^{-rT} | \sigma(T) = \mathsf{s}, \omega] \ \Pr[\sigma(T) = \mathsf{s} | \omega], \end{aligned} \tag{5}$$

the utility for $\sigma \in (\mathsf{s},\mathsf{S})$ is given by

$$\begin{aligned} u_j(\sigma) &= -\frac{c_j}{r} + \Pr[\omega = G] \left(v_j^G + \frac{c_j}{r} \right) \Psi(\sigma,G) + \Pr[\omega = B] \left(v_j^B + \frac{c_j}{r} \right) \Psi(\sigma,B) \\ &+ \Pr[\omega = G] \left(\frac{c_j}{r} \right) \psi(\sigma,G) + \Pr[\omega = B] \left(\frac{c_j}{r} \right) \psi(\sigma,B). \end{aligned}$$

The first line corresponds to the case where the state is good and the upper benchmark S is reached first. The second line is the case where the state is bad but the upper benchmark is reached first, and so on.

As in Stokey (2009), we obtain the following closed-form expressions

$$\begin{array}{ll} \Psi(\sigma,G) = \frac{e^{-R_1(\sigma-\mathsf{s})} - e^{-R_2(\sigma-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} & \Psi(\sigma,B) = \frac{e^{R_2(\sigma-\mathsf{s})} - e^{R_1(\sigma-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})}} \\ \psi(\sigma,B) = \frac{e^{-R_1(\mathsf{S}-\sigma)} - e^{-R_2(\mathsf{S}-\sigma)}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} & \psi(\sigma,G) = \frac{e^{R_2(\mathsf{S}-\sigma)} - e^{R_1(\mathsf{S}-\sigma)}}{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})}} \end{array} \tag{6}$$

with $R_1 = \frac{1}{2}\left(1 - \sqrt{1 + \frac{4r}{\mu'}}\right) < 0$ and $R_2 = \frac{1}{2}\left(1 + \sqrt{1 + \frac{4r}{\mu'}}\right) > 0$, so that $R_1 < R_2$ and $R_1 + R_2 = 1$.

Lemma B0 is key to characterize the shape of the best replies in Lemmas B1 and B2 reported below and proved in Supplementary Appendix B:

Lemma B0 *The conditional probabilities Ψ and ψ satisfy the following properties:*

$$\begin{array}{llll} \textbf{(1)} & \Psi(\sigma,B) = e^{\sigma-\mathsf{S}} \Psi(\sigma,G) & \textbf{(2)} & \psi(\sigma,B) = e^{\sigma-\mathsf{s}} \psi(\sigma,G) \\ \textbf{(3)} & \frac{\partial \Psi(\sigma,G)}{\partial \mathsf{s}} = a \cdot \psi(\sigma,G) < 0 & \textbf{(4)} & \frac{\partial \psi(\sigma,G)}{\partial \mathsf{s}} = b \cdot \psi(\sigma,G) > 0 \\ \textbf{(5)} & \frac{\partial \Psi(\sigma,G)}{\partial \mathsf{S}} = f \cdot \Psi(\sigma,G) < 0 & \textbf{(6)} & \frac{\partial \psi(\sigma,G)}{\partial \mathsf{S}} = g \cdot \Psi(\sigma,G) > 0, \end{array}$$

where

$$\begin{array}{ll} a = \frac{R_1 - R_2}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} < 0 & b = \frac{R_2 e^{R_2(\mathsf{S}-\mathsf{s})} - R_1 e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})}} > 0 \\ f = \frac{R_1 e^{-R_1(\mathsf{S}-\mathsf{s})} - R_2 e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} < 0 & g = \frac{R_2 - R_1}{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})}} > 0. \end{array}$$

Lemma B1 *For a player j with acceptance payoff v_j^G (resp. v_j^B) in the good (resp. bad) state and cost of research c_j per unit of time, for a given s:*

(i) *the upper best reply $B_j(s)$ is independent of q.*

(ii) *$B_j(s) > s$ if $s < \hat{q}_j$ and $B_j(s) = s$ otherwise.*

Lemma B2 *For a player j with acceptance payoff v_j^G (resp. v_j^B) in the good (resp. bad) state and cost of research c_j per unit of time, for a given S:*

(i) the lower best reply $b_j(S)$ is independent of q.

(ii) $b_j(S) < S$ if $S > \hat{q}_j$ and $b_j(S) = S$ otherwise.

Proofs of Main Results

Proof of Proposition 1

(a) This result follows as a special case of Lemma B1 with player j as the informer (with $v_j^G = v_j^B = v_i > 0$ and $c_j = c$). Next, we prove the additional result that for these parameters $B_e(s)$ is decreasing in s for $s < \hat{q}_j$.

Using the log-odds parametrization and applying the implicit function theorem we have

$$\frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}} = -\left.\frac{\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{s} \partial \mathsf{S}}}{\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{s}^2}}\right|_{\mathsf{s}=b_i(\mathsf{S})}.$$

Using (18) in Appendix B, the numerator is given by

$$\left.\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{S} \partial \mathsf{s}}\right|_{\mathsf{s}=b_i(\mathsf{S})} = \frac{e^\sigma}{1+e^\sigma} \psi(\sigma, G) \left(\frac{\partial a}{\partial \mathsf{S}}(1+e^{-\mathsf{S}}) - ae^{-\mathsf{S}}\right) \left(v_i + \frac{c}{r}\right) > 0,$$

and is positive since Lemma B0 established that $a = \frac{R_1 - R_2}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} < 0$ and clearly $\frac{\partial a}{\partial \mathsf{S}} > 0$. The denominator is negative, as shown in (21) in Appendix B. Hence, $b_i(S)$ is increasing in S.

(b) In any period t in which the belief is q, if the informer chooses $\mathscr{A}_i$ (apply for approval), then because of the timing of the informer-authority extensive form game, the game ends after the evaluator's choice. The best reply is then for the evaluator to choose $\mathscr{A}_e$ (approval) if and only if $q \geq \hat{q}_e$ (by definition of $\hat{q}_e$). It follows that, if $q \geq \hat{q}_e$, the unique strategy part of a MPE is for the informer to choose $\mathscr{A}_i$.

For $q < \hat{q}_e$, we have that, if in equilibrium the informer chooses $\mathscr{I}_i$, then the informer must also choose $\mathscr{I}_i$ for any belief q' with $\hat{q}_e > q' > q$. Suppose that this were not the case, and define $\tilde{q}$ the smallest belief $q' \in (q, \hat{q}_e)$, such that $\mathscr{I}_i$ is not chosen at q'. At belief $\tilde{q}$, the informer chooses either $\mathscr{R}_i$, in which case the game ends, or $\mathscr{A}_i$, in which case the evaluator chooses $\mathscr{R}_e$ since $\hat{q}_e > \tilde{q}$. In either case, the outcome at $\tilde{q}$ is R (rejection), so that by choosing $\mathscr{I}_i$ at belief q the informer incurs a cost without a chance of obtaining approval. At belief q the informer would then want to deviate from choosing $\mathscr{I}_i$, reaching a contradiction. This shows that all MPE are characterized by an interval $(\underline{q}, \hat{q}_e)$ where the informer chooses $\mathscr{I}_i$ for $q \in (\underline{q}, \hat{q}_e)$, but not at $q = \underline{q}$.

We now show that $\underline{q} = b_i(\hat{q}_e)$. Suppose first that $\underline{q} > b_i(\hat{q}_e)$. Then, by definition of $b_i(\hat{q}_e)$, the informer has a profitable deviation: at belief $q = \underline{q}$, deviating to $\mathscr{I}_i$ is optimal, a contradiction. Suppose next that $\underline{q} < b_i(\hat{q}_e)$. Then there exists $q'' < b_i(\hat{q}_e)$, such that $q'' \in (\underline{q}, \hat{q}_e)$. By definition of $b_i(\hat{q}_e)$, at $q = q''$ the informer chooses $\mathscr{R}_i$ over $\mathscr{I}_i$ but this is a contradiction since $q'' \in (\underline{q}, \hat{q}_e)$. Thus, in all MPE the informer chooses $\mathscr{I}_i$ for $q \in (b_i(\hat{q}_e), \hat{q}_e)$.

Finally, if $q < b_i(\hat{q}_e)$, the informer is indifferent between $\mathscr{R}_i$ and $\mathscr{A}_i$, given that $\mathscr{A}_i$ will be followed by $\mathscr{R}_e$ and both strategies thus lead to rejection. We conclude that the unique MPE outcome is the one described in Proposition 1.

(c) We derive the payoff achieved by the informer under informer authority $V_i^i(q_0)$ and study its limit when c and r converge to 0 to compare it to the value obtained by KG.[45] First notice that, according to result (b), for all values of c and r, in equilibrium we have $S = \hat{q}_e$ and $s = b_i(\hat{q}_e)$. Using the log-odds parametrization and the characterization of the informer's lower best reply in Lemma B1 (equation 20), we see that, as c converges to 0, s goes to $-\infty$ (when there is no cost of research, the informer never abandons research).

If $q_0 \geq \hat{q}_e$, the outcome of the game under informer authority is therefore A (approval), so that the informer's value is clearly $V_i^i(q_0) = v_i$. By contrast, if $q_0 < \hat{q}_e$, the outcome is I (information acquisition). The utility of the informer in this case is given by expression (14) in Appendix B. Using the fact that under informer authority $\mathsf{S} = \hat{\sigma}_e$, we have that the limit of $V_i^i(\sigma_0)$ as c and r converge to 0 is given by $\frac{e^{\sigma_0}}{1+e^{\sigma_0}} \Psi(\sigma_0, G) \left[v_i (1 + e^{-\hat{\sigma}_e}) \right]$. Substituting $\lim_{c \to 0, r \to 0} R_1 = 0$, $\lim_{c \to 0, r \to 0} R_2 = 1$, and $\lim_{c \to 0, r \to 0} \mathsf{s} = -\infty$ in expression (6) for $\Psi(\sigma, G)$ we see that $\lim_{c \to 0, r \to 0} \Psi(\sigma, G) = 1$. Overall, we find that the limit of $u_i(\sigma_0)$, as c and r go to 0, is $\frac{e^{\sigma_0}}{1+e^{\sigma_0}} \left[v_i (1 + e^{-\hat{\sigma}_e}) \right] = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \frac{1+e^{\hat{\sigma}_e}}{e^{\hat{\sigma}_e}} v_i$, which in the regular space gives $V_i^i(q_0) \to \frac{q_0}{\hat{q}_e} v_i$. The value of the informer is a linear function of the initial belief, equal to the expression derived by KG on pages 2597–2598.

Proof of Proposition 2

(a) This result follows as a special case of Lemma B2 with player j as the evaluator (with $v_j^G = v_e^G$, $v_j^B = v_e^B$ and $c_j = 0$). For this special case we prove the additional result that $B_e(s)$ is decreasing in s for $s < \hat{q}_e$. Using the log-odds parametrization and applying the implicit function theorem, we have

$$\frac{\partial B_e(\mathsf{s})}{\partial \mathsf{s}} = - \left. \frac{\frac{\partial^2 u_e(\sigma)}{\partial \mathsf{S} \partial \mathsf{s}}}{\frac{\partial^2 u_e(\sigma)}{\partial \mathsf{S}^2}} \right|_{\mathsf{S} = B_e(\mathsf{s})} .$$

[45] In the notation for the value function, the subscript the player under consideration (in this case, the informer i) and the superscript refers to the organizational form (in this case, informer authority i).

Expression (15) in the case of the evaluator gives

$$\frac{\partial u_e(\sigma)}{\partial \mathsf{S}} = \frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma,G)\left\{f\cdot\left(v_j^G + e^{-\mathsf{S}}v_j^B\right) - e^{-\mathsf{S}}v_j^B\right\},$$

which implies that the numerator of the above expression is

$$\begin{aligned}\left.\frac{\partial^2 u_e(\sigma)}{\partial \mathsf{S}\partial \mathsf{s}}\right|_{\mathsf{S}=B_e(\mathsf{s})} &= \frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma,G)\frac{\partial f}{\partial \mathsf{s}}\left(v_j^G + e^{-\mathsf{S}}v_j^B\right)\\ &= \frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma,G)\frac{\partial f}{\partial \mathsf{s}}\frac{e^{-\mathsf{S}}v_e^B}{f},\end{aligned}$$

where the second expression is derived from equation (16) that characterizes $B_e(s)$. The numerator is negative given that $f = \frac{R_1 e^{-R_1(\mathsf{S}-\mathsf{s})} - R_2 e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} < 0$ by Lemma B0 and $\frac{\partial f}{\partial \mathsf{s}} = -\frac{(R_1-R_2)^2 e^{-(\mathsf{S}-\mathsf{s})}}{\left(e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}\right)^2} < 0$. The denominator is negative, as shown in (17) in Appendix B. Hence, the evaluator's upper best reply $B_e(\mathsf{s})$ is decreasing in s.

(b) We first show that the outcome for all MPE equilibria of the no-commitment game is characterized by a Wald-cutoff strategy with standards $s^N = b_i(S^N)$ and $S^N = B_e(s^N)$. Define as a maximal interval of research an interval $(\underline{q},\overline{q})$ such that the equilibrium outcome is I (information acquisition) if $q \in (\underline{q},\overline{q})$ and is either A (approval) or R (rejection) at the endpoints $q = \underline{q}$ and for $q = \overline{q}$. As a reminder, we consider Markov Perfect Equilibria where the state is defined by the belief q, so that the equilibrium outcome is fully described by the outcome for each q.

Step 1: *In any Markov perfect equilibrium, the evaluator chooses $\mathscr{R}_e$ over $\mathscr{A}_e$ if and only if $q < \hat{q}_e$.* This follows immediately from the definition of $\hat{q}_e$.

Step 2: *There cannot exist disjoint maximal intervals of research and $\hat{q}_e$ must belong to any maximal interval of research.*

Consider any maximal interval of research $(\underline{q},\overline{q})$ and notice that, by definition, at the endpoints of such an interval the players do not choose strategies $\mathscr{I}_i$ or $\mathscr{I}_e$. At both $q = \underline{q}$ and $q = \overline{q}$, the informer is thus called to decide between $\mathscr{A}_i$ and $\mathscr{R}_i$. If $\underline{q} > \hat{q}_e$, the informer's best reply is to choose $\mathscr{A}_i$, given that the evaluator then chooses $\mathscr{A}_e$, according to Step 1. As a consequence, if $\underline{q} > \hat{q}_e$, the equilibrium outcome is A (approval) at both $q = \underline{q}$ and $q = \overline{q}$. This implies that, at any belief $q \in (\underline{q},\overline{q})$, if the informer chooses $\mathscr{A}_i$, the evaluator chooses $\mathscr{A}_e$. Indeed, $\mathscr{R}_e$ is not a best reply by Step 1 and $\mathscr{I}_e$ also gives a lower payoff since it just delays obtaining the approval payoff. An optimal deviation for the informer at $q \in (\underline{q},\overline{q})$ is thus to choose $\mathscr{A}_i$ with associated outcome A, which contradicts $(\underline{q},\overline{q})$ being a maximal interval of research. We conclude that there cannot exist a maximal interval of research $(\underline{q},\overline{q})$ with $\underline{q} > \hat{q}_e$. A similar argument establishes that we cannot have $\overline{q} < \hat{q}_e$, proving that $\hat{q}_e$ must belong to any maximal interval of research. This directly implies that there cannot exist two disjoint maximal intervals of research.

Step 3: *In all MPE equilibrium the outcome is characterized by a Wald-cutoff strategy with standards $s^N = b_i(S^N)$ and $S^N = B_e(s^N)$.*

Step 2 establishes that in all MPE equilibrium the outcome is characterized by a Wald-cutoff strategy with maximal interval of research $(\underline{q}, \overline{q})$ and that $\hat{q}_e$ must belong to this interval. We now show that $\underline{q}$ and $\overline{q}$ are such that $\underline{q} = b(\overline{q})$ and $\overline{q} = B(\underline{q})$. For any belief $q < \hat{q}_e$ in the interval $(\underline{q}, \overline{q})$, if the informer chooses $\mathscr{A}_i$, the evaluator chooses $\mathscr{I}_e$ (because research is the equilibrium outcome for $q \in (\underline{q}, \overline{q})$). The informer is thus indifferent between choosing $\mathscr{A}_i$ or $\mathscr{I}_i$. At belief $\underline{q}$, by step 1, the outcome is rejection, thus $\underline{q}$ must be a belief such that the informer is indifferent between all three choices and in particular between $\mathscr{I}_i$ and $\mathscr{R}_i$. By definition of b_i this implies that $\underline{q} = b_i(\overline{q})$. If $q > \hat{q}_e$, when the informer chooses $\mathscr{A}_i$, the evaluator chooses between $\mathscr{A}_e$ and $\mathscr{I}_e$. By the same logic and by definition of B_e, it has to be the case that $\overline{q} = B_e(\underline{q})$.

Step 4: *MPE exists and is unique.*

We have established that the equilibrium outcome of the no-commitment game, if it exists, is at the intersection of $b_i(S)$ and $B_e(s)$. We showed in Proposition 1.a that $b_i(S)$ is increasing in S and in part (a) of this proposition that $B_e(s)$ is decreasing in s for $s < \hat{q}_e$ and follows the diagonal for $s \geq \hat{q}_e$. Given that $b_i(S) < S$ for all $S < 1$, $b_i(S)$ is always above the diagonal and thus can cross $B_e(s)$ only once.[46] This implies that, if the MPE exists, it is unique. To see that it exists, we have $b_i(0) = 0$ and $B_e(0) > 0$ and for s large $b_i(S)$ is above the diagonal, so the curves cross once.

(c) Part (b) shows that there is a unique crossing between $b_i(S)$ and $B_e(s)$. Furthermore, $B_e(\hat{q}_e) = \hat{q}_e$ and $b_i(\hat{q}_e) < \hat{q}_e$, so that the two curves necessarily cross at a value $S > \hat{q}_e$. It follows that $S^N > \hat{q}_e$. In turn, since $b_i(S)$ is increasing in S, this implies that $s^N = b_i(S^N) > b_i(\hat{q}_e) = s^i$.

Proof of Proposition 3

The evaluator-commitment outcome solves $\max\limits_{S} u_e|_{s=b_i(S)}$. The solution to this problem, $S^e(q_0)$, depends on the starting belief q_0. Notice that $S^e(q_0)$ is necessarily above the myopic cutoff ($S^e(q_0) \geq \hat{q}_e$) for any q_0, because it is never optimal for the evaluator to approve when the expected payoff is negative. Since $b_i(S)$ is increasing in S, it follows that $b_i(S^e(q_0)) \geq b_i(\hat{q}_e) = s^i$ for any q_0. Given this general property, we distinguish three cases:

(i) $q_0 \in \left(0, s^i\right)$. In this region $q_0 < b_i(S^e(q_0))$, thus the informer does not undertake research.

[46] Note that there is also a crossing for $s = S = 1$, however this equilibrium is not subgame perfect and thus not an MPE since if the informer deviates and submits at $\hat{q}_e < q < 1$, the evaluator would approve.

(ii) $q_0 \in \left(s^i, S^N\right)$. The informer conducts research if $S^e(q_0) < b_i^{-1}(q_0)$, where b_i^{-1} is well defined since the lower best reply of the informer is strictly increasing by Proposition 1.a.ii. Given that the evaluator does not pay the cost of research, for $s^i < q_0 < S^N$, any commitment leading to some research is preferable to no research. Thus, the commitment solution for initial beliefs in this interval is interior. Next, we show that the interior commitment is increasing in q_0. Using the log-odds parametrization and applying the implicit function theorem, we have

$$\frac{\partial S^e}{\partial \sigma_0} = -\left.\frac{\frac{\partial^2 u_e(\sigma_0)}{\partial \mathsf{S} \partial \sigma_0}}{\frac{\partial^2 u_e(\sigma_0)}{\partial \mathsf{S}^2}}\right|_{\mathsf{s}=b_i(\mathsf{S})} \tag{7}$$

Note that

$$\left.\frac{\partial u_e}{\partial \mathsf{S}}\right|_{\mathsf{s}=b_i(\mathsf{S})} = \frac{\partial u_e}{\partial \mathsf{s}}(b_i(\mathsf{S}), \mathsf{S})\ \frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}} + \frac{\partial u_e}{\partial \mathsf{S}}(b_i(\mathsf{S}), \mathsf{S}). \tag{8}$$

Since b_i is independent of σ_0 by Proposition 1.a.i, using expression (8), we have

$$\frac{\partial^2 u_e}{\partial \mathsf{S} \partial \sigma_0}|_{\mathsf{s}=b_i(\mathsf{S})} = \frac{\partial^2 u_e}{\partial \mathsf{s} \partial \sigma_0}(b_i(\mathsf{S}), \mathsf{S})\ \frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}} + \frac{\partial^2 u_e}{\partial \mathsf{S} \partial \sigma_0}(b_i(\mathsf{S}), \mathsf{S}). \tag{9}$$

The first term is positive because its first factor is positive by Lemma C1 and its second factor is also positive by Proposition 1.a.ii. The second term is positive by Lemma C2. Finally, the denominator of (7) is negative by Lemma C3, proving that (7) is positive.[47]

(iii) $q_0 \in \left(S^N, 1\right)$. We show that all interior commitments are such that $S^e(q_0) \leq S_N$. By definition of S^N we have

$$\frac{\partial u_e}{\partial S}\left(b_i(S^N), S^N\right) = \frac{\partial u_e}{\partial S}\left(s^N, S^N\right) = 0. \tag{10}$$

Replacing in expression (8), and using the fact that u_e is decreasing in s and $b_i(S)$ is increasing in S, we obtain $\frac{\partial u_e}{\partial S}|_{s=b_i(S)} < 0$ at $S = S^N$. Thus for $q_0 > S^N$, S^N is the optimal commitment, so that the evaluator chooses to immediately approve for these initial beliefs.

Proof of Proposition 4

(a) Since research is costless, the rejection standard is $s = 0$ for all organizational structures.

(b) According to part (a), $b_i(S) = b_e(S) = 0$ for all $S < 1$, so that we have $S^e = S^N = S_e^*$. Furthermore, for $r > 0$ we have $S_e^* > S_w^*$, given that the evaluator does not internalize the positive externality imposed on the informer upon approval. Thus, $S_w^* \leq S^e = S^N$.

(c) The informer-authority outcome $S^i = \hat{q}_e$ is independent of r; the informer applies for approval as soon as q reaches the myopic cutoff of the evaluator. $S_w^*(r)$ is a strictly decreasing and continuous

[47] See Supplementary Appendix C for Lemmas C1-C3.

function of r, with $S_w^* \to 1$ for $r \to 0$ and $S_w^* \to \hat{q}_w$ for $r \to \infty$. Given that $\hat{q}_w < \hat{q}_e$, there exists $\bar{r}$ such that $S^i \leq S_w^*$ for $r < \bar{r}$ and $S_w^* \leq S^i$ for $r > \bar{r}$.

(d) According to part (c), if $r > \bar{r}$ we have $S_w^* \leq S^i \leq S^e = S^N$ and informer authority unambiguously dominates. For $r < \bar{r}$, S_w^* is decreasing in r and $S^e - S_w^*$ is increasing in r. At the limit when $r \to 0$, we have $S^e - S_w^* \to 0$. So there exists $\hat{r} < \bar{r}$ such that if $r < \hat{r}$ evaluator commitment welfare dominates informer authority and if $r > \hat{r}$ informer authority welfare dominates evaluator commitment.

Proof of Proposition 5

(a) According to Proposition 3, $\hat{q}_e \leq S^e \leq S^N$. Since $b_i(S)$ is strictly increasing in S, we have $b_i(\hat{q}_e) \leq b_i(S^e) \leq b_i(S^N)$. In Proposition 1 we showed $s^i = b_i(\hat{q}_e)$ and by definition, $s^e = b_i(S_e)$ and $s^N = b_i(S^N)$, thus implying $s^i \leq s^e \leq s^N$. Next, we show that $s_w^* \leq s^i$. We have $s_w^* = b_w(S_w^*)$ and $s^i = b_i(\hat{q}_e) = b_w(\hat{q}_e)$ given that $b_i(S)$ and $b_w(S)$ cross at $S = \hat{q}_e$, where the informer imposes no externality on the evaluator. As shown in the main text, $b_w(S)$ is decreasing for $S \in [\hat{q}_w, S_w^*]$ and increasing for $S \in [S_w^*, 1]$. Thus, if $\hat{q}_e \leq S_w^*$, $s^i = b_w(\hat{q}_e)$ is on the decreasing portion of b_w, so that $s_w^* = b_w(S_w^*) \leq s^i = b_w(\hat{q}_e)$. If, instead, $\hat{q}_e > S_w^*$, $s^i = b_w(\hat{q}_e)$ is on the increasing portion of b_w, so that again $s_w^* = b_w(S_w^*) \leq s^i = b_w(\hat{q}_e)$.

(b) By revealed preference, evaluator commitment benefits the evaluator relative to no commitment. The informer also benefits from evaluator commitment given that the standard of approval is decreased ($S^e(q_0) \leq S^N$ for all q_0). Thus, evaluator commitment Pareto dominates no commitment.

(c) $S^i \leq S^e \leq S^N$ was established in Propositions 2 and 3. We now show that $S_w^* \leq S^i$ for either (i) r, (ii) v_i or (iii) c is sufficiently high. When $c \to +\infty$ or $r \to +\infty$, we have $S_w^* \to \hat{q}_w$. Since $\hat{q}_w < \hat{q}_e = S^i$, this establishes that for r or c sufficiently high, $S_w^* \leq S^i$. For result (ii), we use the fact that S_w^* is decreasing in v_i and $S_w^* \to 0$ when $v_i \to +\infty$. Since $S^i = \hat{q}_e$ is independent of v_i, this implies that $S_w^* < S^i$ for v_i large enough.

(d) In accordance with the notation previously introduced, $V_w^e(v_i)$ denotes the social welfare at the evaluator-commitment outcome and $V_w^i(v_i)$ the social welfare at the informer-authority outcome.

(i) Consider first the case $q_0 < \hat{q}_e$. The idea of the proof is the following. We introduce $\underline{v}_i(q_0)$, the value of v_i such that at $q = q_0$ the informer is indifferent between doing research or not under informer authority. If v_i is smaller than $\underline{v}_i(q_0)$, the informer does no research, if it is higher, the informer does research, both under informer authority and evaluator commitment. Since $\underline{v}_i(q_0)$ is decreasing in q_0, we can define $\underline{q}$ such that at $q = \underline{q}$ and $\underline{v}_i(\underline{q})$ informer authority and evaluator commitment both implement the first best. Then if v_i is higher than $\underline{v}_i(\underline{q})$, the planner and the informer have more aligned interest so that informer authority strictly dominates.

Formally, define $\underline{v}_i(q_0)$ the value of v_i such that $b_i(\hat{q}_e) = q_0$. Such a value of v_i exists since $q_0 < \hat{q}_e$, $b_i(\hat{q}_e) = \hat{q}_e$ for $v_i = 0$ and the lower best reply of the informer is strictly decreasing in v_i. According to this definition, if $v_i < \underline{v}_i(q_0)$, then $q_0 < s^i = b_i(\hat{q}_e)$ so that, under informer authority, no research is done and $V_w^i(v_i) = 0$. Similarly, the evaluator under evaluator commitment cannot choose $S \geq \hat{q}_e$ that encourages research by the informer. When $v_i < \underline{v}_i(q_0)$, both organizational forms are equivalent and give a zero payoff to the social planner, i.e., $V_w^i(v_i) = V_w^e(v_i) = 0$.

Define $\check{v}_i$ as the value of v_i at which $S_w^* = \hat{q}_e$. Note that b_i, b_w and B_w all cross at the same point. As v_i increases, b_i shifts to the left while b_w and B_w remain unaffected. Thus, v_i can be chosen such that b_i crosses b_w at the value of S where b_w also crosses B_w. Since b_i and b_w always cross at $S = \hat{q}_e$, the value $\check{v}_i$ such that $S_w^* = \hat{q}_e$ always exists.

Finally, define $\underline{q}$ the value of q_0 such that $\underline{v}_i(\underline{q}) = \check{v}_i$. In this case, for $q_0 = \underline{q}$, if $v_i = \underline{v}_i(q_0)$, informer authority gives the first best. For values of $q_0 < \underline{q}$, if $v_i > \underline{v}_i(q_0)$ (case represented in Figure 6), b_i is to the left of b_w at $q = \hat{q}_e$ so that $S_w^* < \hat{q}_e < S^e$. Informer authority then strictly dominates, given that social welfare is decreasing in S for $S > S_w^*$, moving along $b_w(S)$. Thus, we have characterized the first situation: if $q_0 < \underline{q}$, then for $v_i \leq \underline{v}_i(q_0)$, both organizations are equivalent, while for $v_i > \underline{v}_i(q_0)$, informer authority strictly dominates.[48]

Similarly consider the case $q_0 > \hat{q}_e$. For these initial beliefs, informer authority yields immediate approval with $V_w^i(v_i) = v_i + q_0 v_e^G + (1 - q_0) v_e^B$. The lower best reply of the informer is decreasing in v_i and B_e does not depend on v_i, so that S^N (at the intersection of b_i and B_e) is increasing in v_i. Define $\overline{v}_i(q_0)$ the value of v_i such that $q_0 = S^N$. We have that if $v_i < \overline{v}_i(q_0)$, $q_0 > S^N$ and according to Proposition 3 evaluator commitment leads to immediate approval so that the two organizations are equivalent. In the same spirit as in case 1, only for v_i large does evaluator commitment not lead to immediate approval and for these large values of v_i, not approving immediately in fact leads to lower welfare. Specifically, define $\hat{v}_i$ the value of v_i such that $S^N = S_w^*$ and define $\overline{q}$ the value of q_0 such that $\overline{v}_i(\overline{q}) = \hat{v}_i$. Then, if $q_0 > \overline{q}$, for $v_i \geq \overline{v}_i(q_0)$, $q_0 > S_w^*$ so that the first best is to immediately approve. However, some research is done under evaluator commitment, so that informer authority strictly dominates.

(ii) A direct implication of the definitions of $\check{v}_i$ and $\hat{v}_i$ introduced in (i) is that if $v_i > \tilde{v}_i \equiv max(\check{v}_i, \hat{v}_i)$, informer authority strictly dominates.

[48]Lemma C4 in Appendix C shows situations where $q_0 \in (\underline{q}, \hat{q}_e)$ and evaluator commitment strictly dominates.

(iii) The proof follows the same lines as the proof of (i). Consider first the case $q_0 < \hat{q}_e$. The lower best reply of the informer is increasing in c. As in the proof of (i), define $\underline{v}_i(c)$ as the value of v_i such that $b_i(\hat{q}_e) = q_0$ and $\underline{c}$ as the value of c such that $\underline{v}_i(\underline{c}) = \check{v}_i$, where $S_w^* = \hat{q}_e$. In this case, for values of $c > \underline{c}$, if $v_i > \underline{v}_i(c)$, informer authority strictly dominates. For $q_0 > \hat{q}_e$, can find $\hat{c}$ and $\overline{v}_i(c)$, such that if $c > \hat{c}$ and $v_i \geq \overline{v}_i(c)$, informer authority strictly dominates. Defining $\tilde{c} = \max(\underline{c}, \overline{c})$, we have for $c > \tilde{c}$, informer authority is welfare superior to evaluator commitment and is strictly superior for v_i large.

Proof of Proposition 6

A key property used in this proof is that, under informer authority when the evaluator is a player, the outcome $(s^i(P), S^i = \hat{q}_e(P))$ lies on the planner's lower best reply curve $b_w(S)$ since for $S = \hat{q}_e$ there is no externality, $b_w(S) = b_i(S)$. As P is increased, $\hat{q}_e(P)$ increases and the informer-authority outcome (s^i, S^i) moves up along the planner's lower best reply curve b_w.

(i) $q_0 \in (s^{i_w}, \hat{q}_w)$. In the game with the planner, costly research is undertaken given that $s^{i_w} < q_0$, but this research is socially worthless because $S^{i_w} = \hat{q}_w$; thus the planner obtains a strictly negative payoff by playing directly. If the role of player is delegated to the evaluator, when $q_0 \in (s^{i_w}, s^i]$ informer authority leads to immediate rejection, given that $q_0 \leq s^i$, yielding zero payoff for the planner. As soon as q_0 exceeds s^i, the planner obtains a strictly positive payoff, because the equilibrium outcome under informer authority when the evaluator is a player lies on the planner's lower best reply. Thus, the planner strictly prefers delegating to the evaluator.

(ii) $q_0 \in [\hat{q}_w, S_w^*)$. If the planner plays the game, informer authority results in immediate approval, with a positive payoff given that $q_0 \geq \hat{q}_w$. Define $\overline{S}$ as the value of S that solves $b_w(\overline{S}) = b_w(q_0)$. To see why $\overline{S}$ exists, fix $s = b_w(q_0)$ and consider the planner's value as a function of the initial belief. The smooth-pasting condition is satisfied at $b_w(q_0)$ but not at q_0. For a belief slightly above q_0 the value of the planner is therefore lower than with immediate approval. There exists then an upper threshold $\overline{S}$ such that $b_w(q_0) = b_w(\overline{S})$. Given that the planner is best responding to q_0, when the initial belief is exactly equal to q_0, $\left(s = b_w(q_0), S = \overline{S}\right)$ yields the planner the same payoff as immediate approval.

As P is changed, S^i moves along the planner's lower best reply, as explained at the beginning of the proof. Thus, since $(s = b_w(q_0), S = q_0)$ and $(s = b_w(q_0), S = \overline{S})$ lie on the planner's lower best reply, we can find $\hat{P}(q_0)$ the price level for which $S^i = \overline{S}$ and $\check{P}(q_0)$ the price level for which $S^i = q_0$. When the price level is equal to either $\check{P}(q_0)$ or $\hat{P}(q_0)$, the outcome with delegation lies on the iso-payoff lens corresponding to the approval payoff at q_0, so that the planner is indifferent between delegating and retaining the role of player. For a price in the range between these two levels, $P \in (\check{P}(q_0), \hat{P}(q_0))$, the outcome $\left(s^i, S^i\right)$ of the game between informer and evaluator lies within the lens, as in Figure 8. The planner's payoff at $\left(s^i, S^i\right)$ is thus strictly greater than with immediate approval and the planner strictly benefits from delegating play to the evaluator. If $P < \check{P}(q_0)$, we have that $S^i < q_0$, so that immediate approval is the outcome of the game regardless of whether the evaluator or the planner plays with the informer; thus, the planner is indifferent between delegating and playing. Similarly, when $P > \hat{P}(q_0)$, the outcome $\left(s^i, S^i\right)$ moves up along the social planner's lower reply curve and thus lies outside the lens so that the planner's payoff is now lower than with immediate approval; the planner is thus strictly better off retaining the role of player and approving immediately. This establishes result (ii).

(iii) $q_0 \in [S_w^*, 1)$. The iso-payoff lens corresponding to the approval payoff at q_0 collapses into a point and the interval $(\check{P}(q_0), \hat{P}(q_0))$ becomes empty. Any outcome with research yields the planner a lower payoff than immediate approval. As in case (ii), if $P > \hat{P}(q_0)$, the planner is strictly better off taking up the role of player, but now the planner never strictly benefits from delegation.

Proof of Proposition 7

(**a**) The planner can always block research by committing to an upper cutoff $S^{w_w}(q_0)$ above $b_i^{-1}(q_0)$, inducing the informer to set the lower standard $s^{w_w}(q_0)$ above the initial belief, so that no research is conducted in equilibrium. The incentives of the planner to block are decreasing in q_0. Clearly, for $q_0 \leq s^{w_w}(q_0)$ blocking is optimal since even if the upper threshold was set at S_w^*, the planner would not want to do research. Thus $\tilde{q} \in (s_w^*, \hat{q}_w)$.

(**b**) For intermediate beliefs, $\tilde{q} \leq q_0 < S^{N_w}$, the planner optimally chooses an interior commitment, $S^{w_w}(q_0)$. At $\tilde{q}$, the planner is indifferent between rejecting the project, thus blocking research, and committing to an interior benchmark $S^{w_w}(\tilde{q}) = b_i^{-1}(\tilde{q})$ which lies on the zero-level curve, either way obtaining a zero payoff. If the belief is above $\tilde{q}$, however, an interior commitment becomes strictly preferable, as it induces a pair of standards that are closer than the Nash outcome to the planner's bliss point. How $S^{w_w}(q_0)$ compares to the Nash standard, S^{N_w}, depends on the level of P with three different scenarios. To analyze these cases we define $\check{q} = b_i(\check{S}) = b_w(\check{S})$ as the point of intersection between the lower best replies of the two players. At $\check{q}$ we have $v_i = V(\check{S})$ so that if $S < \check{S}$ the lower best reply of the informer lies below the planner's lower best reply, and conversely if $S > \check{S}$.

(**i**) When $P < P^*$, as in Figure 10, $\tilde{q}$ coincides with the intersection between the planner's lower best reply and the informer's lower best reply, $\tilde{q} = \check{q}$. In this case, if $\tilde{q} \leq q_0 < S^{N_w}$, the planner

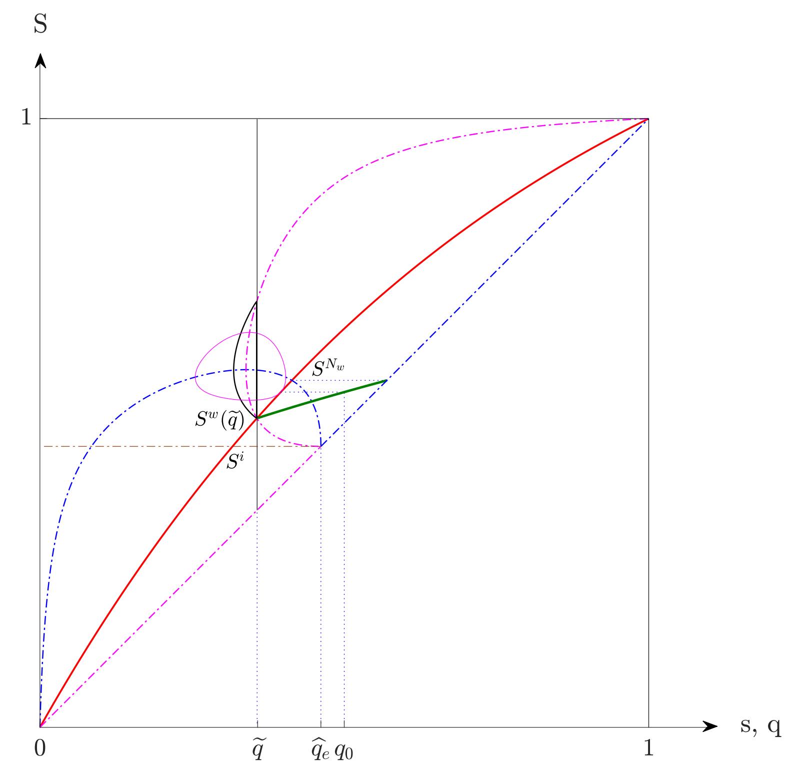

Figure 10: Increasing Commitment Path.

commits to $S^{w_w}(q_0) < S^{N_w}$. Moreover, as $q_0 \to S^{N_w}$, the commitment solution moves along an upward-sloping path that converges to the unique Nash outcome ($S^{w_w}(q_0) \to S^{N_w}$). For a belief greater than $\tilde{q}$, in fact, the iso-payoff curve of the planner (purple in Figure 10) is tangent to the informer's lower best reply for $s^i > \tilde{q}$.

For $P = P^*$, exactly at the boundary between regime (i) and (ii), the lower best replies of the informer and the planner cross exactly at the point where the planner's upper and lower best replies intersect. For this (and only for this) parameter value, the Nash equilibrium coincides with the planner's bliss point, so that commitment has no value and $S^{w_w}(q_0) = S^{N_w}$ for any initial belief $q_0 \in (\tilde{q}, S^{N_w})$.

(ii) When $P^* < P \leq \underline{P}$, as in Figure 11, $\tilde{q}$ still coincides with $\check{q}$. However, the intersection now lies on the upward sloping part of the planner's lower best reply. In this case, if $\tilde{q} \leq q_0 < S_{N_w}$, the optimal commitment is above the Nash level, $S^{w_w}(q_0) > S_{N_w}$. In contrast with the decision-maker commitment, as $q_0 \to S^{N_w}$, the planner's commitment solution moves along a downward (rather than upward) sloping path that converges to the Nash equilibrium of the game, $S^{w_w}(q_0) \to S^{N_w}$. Graphically, the tangency point is now on the left of $\tilde{q}$, as the purple indifference curve in Figure 11 highlights.

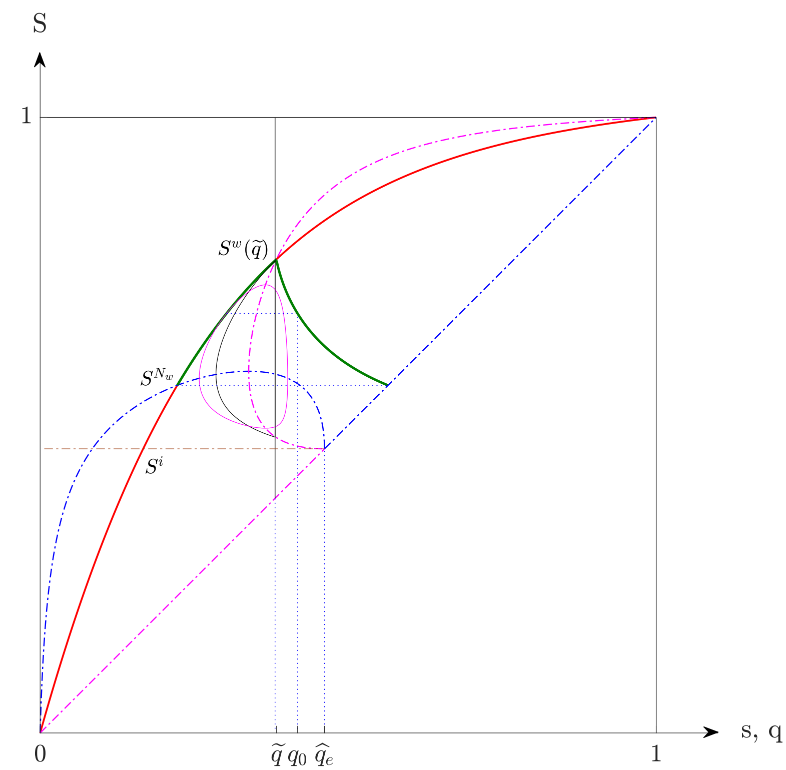

Figure 11: Decreasing Commitment Path.

(iii) For P slightly above $\underline{P}$, the informer's lower best reply (red curve in Figure 9) crosses the planner's zero-level iso-payoff computed at $q_0 = \check{q}$ (dashed black). At this level of the initial belief, thus, the planner chooses an interior commitment, resulting in a strictly positive welfare. There must then be a belief $q_0 = \tilde{q} < \check{q}$ for which the planner obtains exactly zero at the optimal interior commitment; the corresponding iso-payoff (continuous black) is tangent to the informer's lower best reply at a belief, $b_i(S^{w_w}(\tilde{q}))$, which is strictly lower than $\tilde{q}$, as illustrated in Figure 9. At $q_0 = \tilde{q}$ the planner is therefore indifferent between choosing a blocking $S \geq b_i^{-1}(\tilde{q})$ or an interior $S = S^{w_w}(\tilde{q})$ commitment. Since $b_i(S^{w_w}(\tilde{q})) < \tilde{q}$, such an interior commitment induces the informer to set the lower standard below $\tilde{q}$. The planner's optimal interior commitment $S^{w_w}(\tilde{q})$ must then be strictly below the lowest blocking commitment $b_i^{-1}(\tilde{q})$ and this gives rise to the discontinuity displayed in the figure. As in case (ii), as soon as the initial belief increases above $\tilde{q}$, the planner's optimal commitment yields a strictly positive payoff and follows a decreasing path (dark green in Figure 9) which converges to the Nash solution of the game.

Proof of Proposition 8

(a) The result follows directly from the logic of the Coase theorem. If the informer endogenously chooses the price, upon approval request, the informer charges $P(q) = qv_e^G + (1-q)v_e^B$. The preferences of the informer are thus aligned with those of the social planner and the lower best replies are identical. Since the evaluator obtains a zero payoff upon approval regardless of q, under all organizational structures, the first-best outcome (s_w^*, S_w^*) is achieved.

(b) When the informer requests approval, the evaluator sets the price at $P = v_i$. Thus, the informer obtains a zero payoff upon approval, regardless of the belief q at which approval is requested and thus regardless of the organizational form. As a consequence, the informer never performs any research.

Proof of Proposition 9

Consider the subgame following a refusal by evaluator k. In such a subgame, if the informer proposes the technology to evaluator l at belief q, the informer sets price $P(q) = qv_l^G + (1-q)v_l^B - A_l$, in the log-odds space

$$P(\sigma) = \frac{e^\sigma}{1+e^\sigma} v_l^G + \frac{1}{1+e^\sigma} v_l^B - A_l.$$

The utility of the informer at belief σ for an upper benchmark S is then

$$\begin{aligned} u_i(\sigma) &= -\frac{c_i}{r} + \frac{e^\sigma}{1+e^\sigma}\left(P(\mathsf{S}) + \frac{c_i}{r}\right)\Psi(\sigma, G) + \frac{1}{1+e^\sigma}\left(P(\mathsf{S}) + \frac{c_i}{r}\right)\Psi(\sigma, B) \\ &+ \frac{e^\sigma}{1+e^\sigma}\left(\frac{c_i}{r}\right)\psi(\sigma, G) + \frac{1}{1+e^\sigma}\left(\frac{c_i}{r}\right)\psi(\sigma, B). \end{aligned}$$

Using Lemma B0.1 we have

$$\begin{aligned} \frac{e^\sigma}{1+e^\sigma} P(\mathsf{S})\Psi(\sigma, G) + \frac{1}{1+e^\sigma} P(\mathsf{S})\Psi(\sigma, B) &= \frac{e^\sigma}{1+e^\sigma}\Psi(\sigma, G)(1+e^{-\mathsf{S}})P(\mathsf{S}) \\ = \frac{e^\sigma}{1+e^\sigma}\Psi(\sigma, G)(1+e^{-\mathsf{S}})\left(\frac{e^{\mathsf{S}}}{1+e^{\mathsf{S}}} v_l^G + \frac{1}{1+e^{\mathsf{S}}} v_l^B\right) &= \frac{e^\sigma}{1+e^\sigma}\Psi(\sigma, G)\left(v_l^G + e^{-\mathsf{S}} v_l^B\right) \end{aligned}$$

So overall the utility of the informer is equal to

$$u_i(\sigma) = -\frac{c_i}{r} + \frac{e^\sigma}{1+e^\sigma}\Psi(\sigma, G)\left[v_l^G + \mathbf{e}^{-\mathsf{S}} v_l^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_i}{r}\right] + \frac{e^\sigma}{1+e^\sigma}\psi(\sigma, G)(1+e^{-\mathsf{s}})\frac{c_i}{r}.$$

The upper and lower best reply functions are thus identical to those of the evaluator if the evaluator was paying for the research cost and the outcome is denoted by $(\mathsf{s}_l^*, \mathsf{S}_l^*)$, the evaluator's bliss point. This follows the Coasian logic of Proposition 8.

Consider now the offer made to the first evaluator k. If evaluator k refuses the offer, and the state is G (which occurs with probability $\frac{e^\sigma}{1+e^\sigma}$), the payoff of the evaluator is reduced from A_k to D_k at some future uncertain date and with some probability. Building on the definition (5) of the expected discounted probability, denote $\Phi(\mathsf{x},\mathsf{y},\mathsf{z}) := \Psi(\sigma, G)$ for $\sigma = \mathsf{x}$, $\mathsf{S} = \mathsf{y}$ and $\mathsf{S} = \mathsf{z}$, corresponding to the expression in the regular space we use in the main text

$$\phi(x,y,z) := \Phi\left(\log\frac{x}{1-x}, \log\frac{y}{1-y}, \log\frac{z}{1-z}\right). \tag{11}$$

Returning in the log-odds space, if the informer decides to stop at belief σ and producer k refuses, the expected payoff of producer k is

$$A_k - \frac{e^{\sigma}}{1+e^{\sigma}}\Phi(\sigma, \mathsf{s}_l^*, \mathsf{S}_l^*)(A_k - D_k),$$

which according to the definition above and the formula for $\Psi(\sigma, G)$ gives

$$\Phi(\sigma, \mathsf{s}_l^*, \mathsf{S}_l^*) = \frac{e^{-R_1(\sigma - \mathsf{s}_l^*)} - e^{-R_2(\sigma - \mathsf{s}_l^*)}}{e^{-R_1(\mathsf{S}_l^* - \mathsf{s}_l^*)} - e^{-R_2(\mathsf{S}_l^* - \mathsf{s}_l^*)}}.$$

Overall, if the informer decides to stop at belief σ, the informer charges price

$$P(\sigma) = \frac{e^{\sigma}}{1+e^{\sigma}}v_k^G + \frac{1}{1+e^{\sigma}}v_k^B - A_k + \frac{e^{\sigma}}{1+e^{\sigma}}\Phi(\sigma, \mathsf{s}_l^*, \mathsf{S}_l^*)(A_k - D_k).$$

The utility of the informer when stopping at the upper benchmark S and charges $P(\mathsf{S})$ can be written as

$$\begin{aligned} u_i(\sigma) \quad = \quad & -\frac{c_i}{r} + \frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma, G)\left[(v_k^G - A_k) + \Phi(\mathsf{S}, \mathsf{s}_l^*, \mathsf{S}_l^*)(A_k - D_k) + e^{-\mathsf{S}}(v_k^B - A_k) + \left(1 + e^{-\mathsf{S}}\right)\frac{c_i}{r}\right] \\ & + \frac{e^{\sigma}}{1+e^{\sigma}}\psi(\sigma, G)(1 + e^{-\mathsf{s}})\frac{c_i}{r}. \end{aligned}$$

Following the lines of Lemmas B1 and B2, in this case we write the upper and lower best replies by underlying the new terms compared to equations (16 and 19)

$$\begin{aligned} & f\left[(v_k^G - A_k) + \underline{\Phi(\mathsf{S}, \mathsf{s}_l^*, \mathsf{S}_l^*)(A_k - D_k)} + e^{-\mathsf{S}}(v_k^B - A_k) + \left(1 + e^{-\mathsf{S}}\right)\frac{c_i}{r}\right] \\ & + \underline{\frac{\partial\Phi(\mathsf{S}, \mathsf{s}_l^*, \mathsf{S}_l^*)}{\partial \mathsf{S}}(A_k - D_k)} - e^{-\mathsf{S}}\left((v_k^B - A_k) + \frac{c_i}{r}\right) + g.(1 + e^{-\mathsf{s}})\frac{c_i}{r} = 0, \end{aligned} \tag{12}$$

$$\begin{aligned} & (v_k^G - A_k) + \underline{\Phi(\mathsf{S}, \mathsf{s}_l^*, \mathsf{S}_l^*)(A_k - D_k)} + e^{-\mathsf{S}}(v_k^B - A_k) + \left(1 + e^{-\mathsf{S}}\right)\frac{c_i}{r} \\ = \quad & -\frac{1}{a}\frac{c_i}{r}\left[b(1 + e^{-\mathsf{s}}) - e^{-\mathsf{s}}\right]. \end{aligned} \tag{13}$$

We denote this outcome $(\tilde{s}_k, \tilde{S}_k)$. Using the above observations, we now establish the results of Proposition 9.

(a) Direct inspection of the lower best reply in (13) shows that an increase in the externality (a decrease in D_k) is equivalent to an increase in v_k^G, and thus pushes down the lower best reply. Indeed an increase in the payoff in the good state encourages the informer to do more research at the bottom.

For the upper best reply (12), there are two countervailing effects. An increase in the externality increases the payoff of the informer in the good state and thus makes delaying acceptance costly. However, delaying acceptance means that upon acceptance $\Phi(\tilde{\mathsf{S}}_k, \mathsf{s}_l^*, \mathsf{S}_l^*)$ is larger since $\tilde{\mathsf{S}}_k$ is closer to l's approval benchmark S_l^* and further from the rejection benchmark s_l^*; thus, there is a higher externality that can be extracted. We have

$$\frac{\partial B_i(\mathsf{s})}{\partial D_k} = -\left.\frac{\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{S} \partial D_k}}{\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{S}^2}}\right|_{\tilde{\mathsf{S}}_k = B_i(\tilde{\mathsf{s}}_k)},$$

where

$$\begin{aligned}
\left.\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{S} \partial D_k}\right|_{\tilde{\mathsf{S}}_k = B_i(\tilde{\mathsf{s}}_k)} &= -\frac{e^\sigma}{1+e^\sigma}\Psi(\sigma, G)\left(f\Phi(\tilde{\mathsf{S}}_k, \mathsf{s}_l^*, \mathsf{S}_l^*) + \frac{\partial \Phi(\tilde{\mathsf{S}}_k, \mathsf{s}_l^*, \mathsf{S}_l^*)}{\partial \mathsf{S}}\right) \\
&= -\frac{e^\sigma}{1+e^\sigma}\Psi(\sigma, G)\left(f\Phi(\tilde{\mathsf{S}}_k, \mathsf{s}_l^*, \mathsf{S}_l^*) - f_2\Phi(\tilde{k}_i, \mathsf{s}_l^*, \mathsf{S}_l^*)\right)
\end{aligned}$$

with

$$f = \frac{R_1 e^{-R_1(\tilde{\mathsf{S}}_k - \tilde{\mathsf{s}}_k)} - R_2 e^{-R_2(\tilde{\mathsf{S}}_k - \tilde{\mathsf{s}}_k)}}{e^{-R_1(\tilde{\mathsf{S}}_k - \tilde{\mathsf{s}}_k)} - e^{-R_2(\tilde{\mathsf{S}}_k - \tilde{\mathsf{s}}_k)}} \text{ and } f_2 = \frac{R_1 e^{-R_1(\tilde{\mathsf{S}}_k - \mathsf{s}_l^*)} - R_2 e^{-R_2(\tilde{\mathsf{S}}_k - \mathsf{s}_l^*)}}{e^{-R_1(\tilde{\mathsf{S}}_k - \mathsf{s}_l^*)} - e^{-R_2(\tilde{\mathsf{S}}_k - \mathsf{s}_l^*)}}.$$

Defining $\varphi(x) = \frac{R_1 e^{-R_1 x} - R_2 e^{-R_2 x}}{e^{-R_1 x} - e^{-R_2 x}}$, we now show that $\varphi(x)$ is increasing in x. We have $\tilde{\mathsf{S}}_k < \mathsf{S}_l^*$ and furthermore, necessarily in equilibrium $v_k^G + \Phi(\mathsf{S}, \mathsf{s}_l^*, \mathsf{S}_l^*)(A_k - D_k) > v_l^G$. If not, the informer would have offered the project to evaluator l. This implies that $\tilde{\mathsf{s}}_k < \mathsf{s}_l^*$, so that $f = \varphi(\tilde{\mathsf{s}}_k) < f_2 = \varphi(\mathsf{s}_l^*)$. We thus conclude that $\left.\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{S} \partial D_k}\right|_{\tilde{\mathsf{S}}_k = B_i(\tilde{\mathsf{s}}_k)} > 0$, so that $\frac{\partial B_i(\mathsf{s})}{\partial D_k} > 0$, meaning that the upper best reply is decreasing in the externality. So overall an increase in the externality decreases $\tilde{\mathsf{S}}_i$ and $\tilde{\mathsf{s}}_i$.

(b) Evaluator k at beliefs σ obtains payoff

$$A_k - \frac{e^\sigma}{1+e^\sigma}\,\Psi(\sigma, G)\,\Phi(\tilde{\mathsf{S}}_k, \mathsf{s}_l^*, \mathsf{S}_l^*)(A_k - D_k).$$

The first evaluator is clearly worse off. The second evaluator incurs the externality with some probability and thus is also worse off.

(c) Denote $D = D_k = D_l$ and $A = A_k = A_l$ the common values of profits for the two evaluators k and l. Since $v_1^G > v_2^G$, $\mathsf{s}_1^* < \mathsf{s}_2^*$ and $\mathsf{S}_1^* < \mathsf{S}_2^*$, for any S, $\Phi(\mathsf{S}, \mathsf{s}_1^*, \mathsf{S}_1^*) < \Phi(\mathsf{S}, \mathsf{s}_2^*, \mathsf{S}_2^*)$. Hence, when the externality $A - D$ is sufficiently large we have, for any S,

$$v_1^G + \Phi(\mathsf{S}, \mathsf{s}_2^*, \mathsf{S}_2^*)(A - D) < v_2^G + \Phi(\mathsf{S}, \mathsf{s}_1^*, \mathsf{S}_1^*)(A - D).$$

Thus the informer offers the technology to the least efficient evaluator, regardless of the choice of S determining when research is stopped.

Proof of Proposition 10

As derived in the main text, we have that the entrant accepts the project only for beliefs above its myopic cutoff $\hat{q}_E = \frac{P}{\Pi_E}$. On the other hand, the incumbent , when the first to be approached, knows that he will suffer an externality if the state is G and the myopic cutoff of the entrant $\hat{q}_E$ is reached. In case of rejection, the informer thus researches in $(b_i(\hat{q}_E), \hat{q}_E)$. If approached first, the incumbent thus accepts the project for beliefs above the competitive cutoff defined by

$$\tilde{q}_I = \frac{P}{\Pi_{m1} - \Pi_{m0} + \phi(\tilde{q}_I, b_i(\hat{q}_E), \hat{q}_E)(\Pi_{m0} - \Pi_I)}.$$

If $\Pi_E < \Pi_{m1} - \Pi_{m0}$, regardless of the value of $\phi(\tilde{q}_I, b_i(\hat{q}_E), \hat{q}_E) \in (0,1)$, we have $\hat{q}_E > \tilde{q}_I$ and the informer optimally chooses to approach the incumbent first. If, instead, $\Pi_E > \Pi_{m1} - \Pi_{m0}$, we see that if $\phi(\tilde{q}_I, b_i(\hat{q}_E), \hat{q}_E) \to 0$, the entrant will be approached first since $\Pi_E > \Pi_{m1} - \Pi_{m0}$. On the other hand if $\phi(\tilde{q}_I, b_i(\hat{q}_E), \hat{q}_E) \to 1$, the incumbent will be approached first since $\Pi_E < \Pi_{m1} - \Pi_I$. Finally we have that $\tilde{q}_I$ is decreasing in ϕ. Given that for any (q_1, q_2, q_3), $\phi(q_1, q_2, q_3)$ is decreasing in μ, we can find a value $\mu^*(c)$, such that the informer approaches the incumbent first if and only if $\mu > \mu^*$.

Parameters Used in Figures. Figure 1: $\mu = 12$, $v_i = 1.7$, $c = 15$, $r = 5$, $v_e^G = 0.5$, $v_e^G = -0.5$, $q = 0.45$. Figure 2: Solid line: $v_i = 1$, $c = 0$, $r = 0$; dashed curve: $v_i = 1$, $c = 0$, $r = 3$; dotted curve: $v_i = 1$, $c = 7$, $r = 4$. Figure 3: $\mu = 12$, $v_i = 1.7$, $c = 15$, $r = 5$, $v_e^G = 0.5$, $v_e^G = -0.5$, $q = 0.45$. Figure 4: $\mu = 12$, $v_i = 1.7$, $c = 15$, $r = 5$, $v_e^G = 0.5$, $v_e^G = -0.5$, $q = 0.45$. Figure 5: $\mu = 5$, $v_i = 2.5$, $v_e^G = 6$, $v_e^B = -6$, $c = 0$, $r = 2.5$, $q = 0.3$. Figure 6: $\mu = 5$, $v_i = 2$, $v_e^G = 4$, $v_e^B = -6$, $c = 15$, $r = 0.8$. Figure 7: $\mu = 5$, $v_i = 0.36$, $v_e^G = 4$, $v_e^B = -6$, $c = 15$, $r = 0.8$. Figure 8: $\mu = 5$, $v_i = 0.6$, $v_e^G = 6$, $v_e^B = -6$, $c = 15$, $r = 0.8$, $P = 1.5$. Figure 9: $\mu = 5$, $v_i = 0.5$ $v_e^G = 7$, $v_e^B = -6$, $c = 15$, $r = 0.8$, $P = 4.7$. Figure 10: $\mu = 5$, $v_i = 0.5$, $v_e^G = 7$, $v_e^B = -6$, $c = 15$, $r = 0.8$, $P = 0.1$. Figure 11: $\mu = 5$, $v_i = 0.5$, $v_e^G = 7$, $v_e^G = -6$, $c = 15$, $r = 0.8$, $P = 3.5$.

B Supplementary Appendix B: Wald Benchmark Proofs

Proof of Lemma B0

A direct computation yields the following expressions for the conditional probabilities

$$\begin{aligned}\Psi(\sigma,B) &= \frac{e^{R_2(\sigma-\mathsf{s})}-e^{R_1(\sigma-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}=\frac{1}{e^{(\mathsf{S}-\mathsf{s})}}\frac{e^{R_2(\sigma-\mathsf{s})}-e^{R_1(\sigma-\mathsf{s})}}{e^{(R_2-1)(\mathsf{S}-\mathsf{s})}-e^{(R_1-1)(\mathsf{S}-\mathsf{s})}}\\ &= e^{-(S-s)}\frac{e^{R_2(\sigma-\mathsf{s})}-e^{R_1(\sigma-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}=e^{\sigma-S}\frac{e^{-R_1(\sigma-\mathsf{s})}-e^{-R_2(\sigma-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}=e^{\sigma-S}\Psi(\sigma,G)\end{aligned}$$

and

$$\begin{aligned}\psi(\sigma,B) &= \frac{e^{-(1-R_2)(\mathsf{S}-\sigma)}-e^{-(1-R_1)(\mathsf{S}-\sigma)}}{e^{-(1-R_2)(\mathsf{S}-\mathsf{s})}-e^{-(1-R_1)(\mathsf{S}-\mathsf{s})}}=\frac{e^{-\mathsf{S}+\sigma+R_2(\mathsf{S}-\sigma)}-e^{-\mathsf{S}+\sigma+R_1(\mathsf{S}-\sigma)}}{e^{-\mathsf{S}+\mathsf{s}+R_2(\mathsf{S}-\mathsf{s})}-e^{-\mathsf{S}+\mathsf{s}+R_1(\mathsf{S}-\mathsf{s})}}\\ &= \frac{e^{\sigma-\mathsf{S}}}{e^{-(\mathsf{S}-\mathsf{s})}}\frac{e^{R_2(\mathsf{S}-\sigma)}-e^{R_1(\mathsf{S}-\sigma)}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}=e^{\sigma-\mathsf{s}}\psi(\sigma,G).\end{aligned}$$

This establishes parts (1) and (2) of Lemma B0.

Taking the derivative of $\Psi(\sigma,G)$ with respect to s and rearranging terms we obtain

$$\begin{aligned}\frac{\partial\Psi(\sigma,G)}{\partial\mathsf{s}} &= (R_1-R_2)\frac{e^{-R_1(\mathsf{S}-\mathsf{s})-R_2(\sigma-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})-R_1(\sigma-\mathsf{s})}}{\left(e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}\right)^2}=(R_1-R_2)e^{\mathsf{s}-\sigma}\frac{e^{-R_1(\mathsf{S}-\sigma)}-e^{-R_2(\mathsf{S}-\sigma)}}{\left(e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}\right)^2}\\ &= \frac{(R_1-R_2)e^{\mathsf{s}-\sigma}\psi(\sigma,B)}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}=\frac{(R_1-R_2)\psi(\sigma,G)}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}=a\psi(\sigma,G),\end{aligned}$$

where $a<0$, since $e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}>0$ and $R_1-R_2<0$, and a is independent of σ. Similarly, for $\psi(\sigma,G)$ we have

$$\begin{aligned}\frac{\partial\psi(\sigma,G)}{\partial\mathsf{s}} &= -\frac{\left(-R_2e^{R_2(\mathsf{S}-\mathsf{s})}+R_1e^{R_1(\mathsf{S}-\mathsf{s})}\right)\left(e^{R_2(\mathsf{S}-\sigma)}-e^{R_1(\mathsf{S}-\sigma)}\right)}{\left(e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}\right)^2}\\ &= \frac{R_2e^{R_2(\mathsf{S}-\mathsf{s})}-R_1e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}\psi(\sigma,G)=b\psi(\sigma,G).\end{aligned}$$

where $b>0$, since both $e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}>0$ and $R_2e^{R_2(\mathsf{S}-\mathsf{s})}-R_1e^{R_1(\mathsf{S}-\mathsf{s})}>0$, and b is independent of σ. This proves parts (3) and (4).

Finally, taking the derivative of $\Psi(\sigma,G)$ with respect to S we obtain

$$\begin{aligned}\frac{\partial\Psi(\sigma,G)}{\partial S} &= -\frac{\left(e^{-R_1(\sigma-\mathsf{s})}-e^{-R_2(\sigma-\mathsf{s})}\right)\left(-R_1e^{-R_1(\mathsf{S}-\mathsf{s})}+R_2e^{-R_2(\mathsf{S}-\mathsf{s})}\right)}{\left(e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}\right)^2}\\ &= \frac{R_1e^{-R_1(\mathsf{S}-\mathsf{s})}-R_2e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}\Psi(\sigma,G)=f\Psi(\sigma,G),\end{aligned}$$

where $f<0$, since $e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}>0$ and $R_1e^{-R_1(\mathsf{S}-\mathsf{s})}<0<R_2e^{-R_2(\mathsf{S}-\mathsf{s})}$, and f is independent of σ. Similarly, we have

$$\begin{aligned}\frac{\partial\psi(\sigma,G)}{\partial\mathsf{S}} &= (R_2-R_1)\frac{e^{R_1(\mathsf{S}-\sigma)+R_2(\mathsf{S}-\mathsf{s})}-e^{R_2(\mathsf{S}-\sigma)+R_1(\mathsf{S}-\mathsf{s})}}{\left(e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}\right)^2}=(R_2-R_1)\frac{e^{\mathsf{S}-\sigma}\left(e^{R_2(\sigma-\mathsf{s})}-e^{R_1(\sigma-\mathsf{s})}\right)}{\left(e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}\right)^2}\\ &= \frac{(R_2-R_1)\,e^{\mathsf{S}-\sigma}\Psi(\sigma,B)}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}=\frac{(R_2-R_1)\,\Psi(\sigma,G)}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}=g\Psi(\sigma,G)>0.\end{aligned}$$

where $g>0$, since both $R_2-R_1>0$ and $e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}>0$, and g does not depend on σ. This completes the proof of Lemma B0.

Proof of Lemma B1

We provide the most general characterization for the upper best reply $B_j(\mathsf{s})$ for a player j who gets a payoff v_j^G (v_j^B) in the good (bad) state and pays a cost of research c_j per unit of time.

(i) We examine in turn the first-order then second-order conditions and show that the first order condition is independent of σ.

First-Order Condition for the Upper Best Reply

By parts (1) and (2) of Lemma B0 the utility function of player j, $u_j(\sigma)$, can be written as

$$u_j(\sigma)=-\frac{c_j}{r}+\frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma,G)\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]+\frac{e^{\sigma}}{1+e^{\sigma}}\psi(\sigma,G)(1+e^{-\mathsf{s}})\frac{c_j}{r}. \quad (14)$$

By parts (5) and (6) of Lemma B0**,** taking the derivative with respect to S then yields

$$\frac{\partial u_j(\sigma)}{\partial\mathsf{S}}=\frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma,G)\left\{f\cdot\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]-e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)+g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}\right\}, \quad (15)$$

which implies that, at an interior solution, the following first-order condition must be satisfied

$$f\cdot\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]=e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)-g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}. \quad (16)$$

Equation (16) establishes that $B_j(\mathsf{s})$ is independent of σ in the log-odds space, or, equivalently, that $B_j(s)$ is independent of q in the regular space. Furthermore, it implies that $v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}>0$ must hold at $\mathsf{S}=B_i(\mathsf{s})$. Two cases can, in fact, be distinguished: if $e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)\geq 0$, then $v_j^G+e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)+\frac{c_j}{r}>0$ simply follows from $v_j^G>0$ and $\frac{c_j}{r}>0$. If $e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)<0$, then $f\left[v_j^G+e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)+\frac{c_j}{r}\right]<0$ must hold, since $g\cdot(1+e^{-\mathsf{s}})>0$ and $f<0$, so that $v_j^G+e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)+\frac{c_j}{r}>0$ is again satisfied.

In the case of the evaluator, where $c_e = 0$, (16) simplifies into

$$v_e^G + e^{-\mathsf{S}} v_e^B = \frac{e^{-\mathsf{S}} v_e^B}{f}.$$

Note that in the regular space the first order condition can be expressed as equation (2) in the main text

$$-\beta_3(s,S)\, v_e^B = \beta_4(s,S)\, v_A(S),$$

where

$$\begin{aligned}
\beta_3(s,S) &= (1-S) \\
\beta_4(s,S) &= -f(s,S) > 0.
\end{aligned}$$

Second-Order Condition for the Upper Best Reply

Taking a derivative with resect to S of (15) yields

$$\begin{aligned}
&\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2} \\
= \;& \frac{e^{\sigma}}{1+e^{\sigma}} \left\{ \begin{array}{c} \frac{\partial \Psi(\sigma,G)}{\partial \mathsf{S}} \{ f \cdot [v_j^G + e^{-\mathsf{S}} v_j^B + (1-e^{-S}) \frac{c_j}{r}] - e^{-\mathsf{S}} (v_j^B + \frac{c_j}{r}) + g \cdot (1+e^{-\mathsf{s}}) \frac{c_j}{r} \} \\ +\Psi(\sigma,G) \left\{ \frac{\partial f}{\partial \mathsf{S}} [v_j^G + e^{-\mathsf{S}} v_j^B + (1+e^{-S}) \frac{c_j}{r}] + e^{-\mathsf{S}} (v_j^B + \frac{c_j}{r})(1-f) + \frac{\partial g}{\partial \mathsf{S}} (1+e^{-\mathsf{s}}) \frac{c_j}{r} \right\} \end{array} \right\}.
\end{aligned}$$

Equation (16) then implies

$$\begin{aligned}
&\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_j(\mathsf{s})} \\
= \;& \frac{e^{\sigma}}{1+e^{\sigma}} \Psi(\sigma,G) \left\{ e^{-\mathsf{S}} \left(v_j^B + \frac{c_j}{r} \right) \left[\frac{\partial f}{\partial \mathsf{S}} \frac{1}{f} + (1-f) \right] + \left(\frac{\partial g}{\partial \mathsf{S}} - \frac{\partial f}{\partial \mathsf{S}} \frac{g}{f} \right) (1+e^{-\mathsf{s}}) \frac{c_j}{r} \right\} \\
= \;& \frac{e^{\sigma}}{1+e^{\sigma}} \Psi(\sigma,G) \left\{ e^{-\mathsf{S}} \left(v_j^B + \frac{c_j}{r} \right) \left[\frac{\partial f}{\partial \mathsf{S}} \frac{1}{f} + (1-f) \right] + g \cdot \left(\frac{\partial g}{\partial \mathsf{S}} \frac{1}{g} - \frac{\partial f}{\partial \mathsf{S}} \frac{1}{f} \right) (1+e^{-\mathsf{s}}) \frac{c_j}{r} \right\}.
\end{aligned}$$

Some algebra yields

$$\begin{aligned}
1-f &= \frac{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})} - R_1 e^{-R_1(\mathsf{S}-\mathsf{s})} + R_2 e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} \\
&= \frac{R_2 e^{-R_1(\mathsf{S}-\mathsf{s})} - R_1 e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} = \frac{R_2 e^{R_2(\mathsf{S}-\mathsf{s})} - R_1 e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})}} = -\frac{\partial g}{\partial \mathsf{S}} \frac{1}{g}.
\end{aligned}$$

Substituting for $\frac{\partial g}{\partial \mathsf{S}} \frac{1}{g}$ in the above expression and rearranging terms we have

$$\begin{aligned}
&\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_i(\mathsf{s})} \\
= \;& \frac{e^{\sigma}}{1+e^{\sigma}} \Psi(\sigma,G) \left\{ e^{-\mathsf{S}} \left(v_j^B + \frac{c_j}{r} \right) \left[\frac{\partial f}{\partial \mathsf{S}} \frac{1}{f} + (1-f) \right] + g \left[-(1-f) - \frac{\partial f}{\partial \mathsf{S}} \frac{1}{f} \right] (1+e^{-\mathsf{s}}) \frac{c_j}{r} \right\} \\
= \;& \frac{e^{\sigma}}{1+e^{\sigma}} \Psi(\sigma,G) \left[\frac{\partial f}{\partial \mathsf{S}} \frac{1}{f} + (1-f) \right] \left[e^{-\mathsf{S}} \left(v_j^B + \frac{c_j}{r} \right) - g \cdot (1+e^{-\mathsf{s}}) \frac{c_j}{r} \right]
\end{aligned}$$

which, by equation (16), can be rewritten as

$$\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_j(\mathsf{s})} = \frac{e^\sigma}{1+e^\sigma}\Psi(\sigma,G)\left[\frac{\partial f}{\partial \mathsf{S}} + f\cdot(1-f)\right]\left[v_j^G + e^{-\mathsf{S}}v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right].$$

Recalling from above that $v_j^G + e^{-\mathsf{S}}v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r} > 0$ at $\mathsf{S} = B_j(\mathsf{s})$, one sees that $\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_j(\mathsf{s})}$ is negative if and only if $\frac{\partial f}{\partial \mathsf{S}} + f\cdot(1-f) < 0$, that is if and only if $\frac{\partial f}{\partial \mathsf{S}} < -f(1-f)$. We therefore show that the last inequality is indeed satisfied. A direct calculation gives

$$\frac{(R_2-R_1)^2 e^{-(\mathsf{S}-\mathsf{s})}}{\left(e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}\right)^2} < \frac{\left(R_2^2+R_1^2\right)e^{-(\mathsf{S}-\mathsf{s})} - R_1R_2\left(e^{-2R_1(\mathsf{S}-\mathsf{s})} + e^{-2R_2(\mathsf{S}-\mathsf{s})}\right)}{\left(e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}\right)^2},$$

which holds if and only if

$$2e^{-(\mathsf{S}-\mathsf{s})} < e^{-2R_1(\mathsf{S}-\mathsf{s})} + e^{-2R_2(\mathsf{S}-\mathsf{s})}.$$

Given the expressions for R_1 and R_2, this inequality is satisfied, as the function $e^{-(1-x)(\mathsf{S}-\mathsf{s})} + e^{-(1+x)(\mathsf{S}-\mathsf{s})}$ is increasing in x. Since $\frac{\partial f}{\partial \mathsf{S}} + f\cdot(1-f) < 0$ and $v_j^G + e^{-\mathsf{S}}v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r} > 0$, we can finally conclude that

$$\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_j(\mathsf{s})} < 0. \tag{17}$$

(ii) We now examine the slope of the upper best reply. First, we show that $B_j(\mathsf{s}) > \mathsf{s}$ if $\mathsf{s} < \hat{\sigma}_j$ and $B_j(\mathsf{s}) = \mathsf{s}$ otherwise. We start with computing the limit of $\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}$ as $\mathsf{S} \to \mathsf{s}$. Recall that

$$\frac{\partial u_j(\sigma)}{\partial \mathsf{S}} = \frac{e^\sigma}{1+e^\sigma}\Psi(\sigma,G)\left\{f\cdot\left[v_j^G + e^{-\mathsf{S}}v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right] - e^{-\mathsf{S}}\left(v_j^B + \frac{c_j}{r}\right) + g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}\right\}$$

and focus on the last term of the product. A simple calculation gives

$$\begin{aligned} &\lim_{\mathsf{S}\to\mathsf{s}}\left\{f\cdot\left[v_j^G + e^{-\mathsf{S}}v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right] - e^{-\mathsf{S}}\left(v_j^B + \frac{c_j}{r}\right) + g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}\right\} \\ = \quad &\lim_{\mathsf{S}\to\mathsf{s}} f\cdot\left[v_j^G + e^{-\mathsf{s}}v_j^B\right] - e^{-\mathsf{s}}\left(v_j^B + \frac{c_j}{r}\right) + \lim_{\mathsf{S}\to\mathsf{s}}(f+g)\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}. \end{aligned}$$

Because $\lim_{\mathsf{S}\to\mathsf{s}} f = -\infty$ and $\lim_{\mathsf{S}\to\mathsf{s}}(f+g) = 0$, one sees that the sign of the limit above depends on the sign of $v_j^G + e^{-\mathsf{s}}v_j^B$. Specifically, we have

$$\lim_{\mathsf{S}\to\mathsf{s}}\left\{f\cdot\left[v_j^G + e^{-\mathsf{S}}v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right] - e^{-\mathsf{S}}\left(v_j^B + \frac{c_j}{r}\right) + g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}\right\} = \infty$$

if $\mathsf{s} < \hat{\sigma}_j$, in which case $v_j^G + e^{-\mathsf{s}}v_j^B < 0$, and

$$\lim_{\mathsf{S}\to\mathsf{s}}\left\{f\cdot\left[v_j^G + e^{-\mathsf{S}}v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right] - e^{-\mathsf{S}}\left(v_j^B + \frac{c_j}{r}\right) + g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}\right\} = -\infty$$

otherwise.

Since $\lim_{\mathsf{S}\to\mathsf{s}} \frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma,G)=\infty$, overall we have $\lim_{\mathsf{S}\to\mathsf{s}} \frac{\partial u_j(\sigma)}{\partial \mathsf{S}}=\infty$ if $\mathsf{s}<\hat{\sigma}_j$ and $\lim_{\mathsf{S}\to\mathsf{s}} \frac{\partial u_j(\sigma)}{\partial \mathsf{S}}=-\infty$ if $\mathsf{s}\geq\hat{\sigma}_j$.

Next, we compute the limit of $\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}$ as $\mathsf{S}\to\infty$. We have

$$\begin{aligned}&\lim_{\mathsf{S}\to\infty}\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}\\ =\;&\lim_{\mathsf{S}\to\infty}\frac{e^{\sigma}\Psi(\sigma,G)}{1+e^{\sigma}}\left\{f\cdot\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]-e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)+g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}\right\}.\end{aligned}$$

Focusing on the second term of the product, we obtain

$$\begin{aligned}&\lim_{\mathsf{S}\to\infty}\left\{f\cdot\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]-e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)+g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}\right\}\\ =\;&\lim_{\mathsf{S}\to\infty}f\cdot\left[v_j^G+\frac{c_j}{r}\right]+\lim_{\mathsf{S}\to\infty}g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}.\end{aligned}$$

Since $\lim_{\mathsf{S}\to\infty}\frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma,G)=0$, $\lim_{\mathsf{S}\to\infty}f=R_1<0$ and $\lim_{\mathsf{S}\to\infty}g=0$, we have that overall $\lim_{\mathsf{S}\to\infty}\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}=0^-$.

Having computed the limits at the two extremes of the domain of S, we now consider two different cases. First, assume $\mathsf{s}<\hat{\sigma}_e$. Then, since $\lim_{\mathsf{S}\to\mathsf{s}}\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}=\infty$ and $\lim_{\mathsf{S}\to\infty}\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}=0^-$, by continuity there must exist a solution to $\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}=0$, implying that in this case $B_j(\mathsf{s})>\mathsf{s}$. Next, suppose $\mathsf{s}\geq\hat{\sigma}_j$. In this case we show that $\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}<0$. To see this assume by contradiction that there exists s_1 such that $\left.\frac{\partial u(\sigma)}{\partial \mathsf{S}}\right|_{\mathsf{S}=B_j(\mathsf{s}_1)}\geq 0$. Two cases can then be distinguished. If $\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_j(\mathsf{s}_1)}\geq 0$, we have a contradiction since we have shown $\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_j(\mathsf{s})}<0$ for any s. Otherwise, since $\lim_{\mathsf{S}\to\mathsf{s}}\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}=-\infty$ and $\lim_{\mathsf{S}\to\infty}\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}=0^-$, there must exist s_2 such that $\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_j(\mathsf{s}_2)}\geq 0$, which again yields a contradiction. This establishes that $B_j(\mathsf{s})>\mathsf{s}$ if $\mathsf{s}<\hat{\sigma}_j$ and $B_j(\mathsf{s})=\mathsf{s}$ otherwise.

Proof of Lemma B2

We provide the most general characterization for the lower best reply $b_j(\mathsf{S})$ for a player j who gets a payoff v_j^G (v_j^B) in the good (bad) state and pays a cost of research c_j per unit of time.

(i) We examine in turn the first-order then second-order conditions and show that the first order condition is independent of σ.

First-Order Condition for the Lower Best Reply

By parts (3) and (4) of Lemma B0**,** taking a derivative of (14) with respect to s yields

$$\frac{\partial u_j(\sigma)}{\partial \mathsf{s}}=\frac{e^{\sigma}}{1+e^{\sigma}}\psi(\sigma,G)\left\{a\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]+\frac{c_j}{r}\left[b(1+e^{-\mathsf{s}})-e^{-\mathsf{s}}\right]\right\}. \tag{18}$$

Hence, player j's first order condition is

$$v_j^G + e^{-\mathsf{S}} v_j^B + \left(1 + e^{-\mathsf{S}}\right) \frac{c_j}{r} = -\frac{1}{a}\frac{c_j}{r}\left[b(1+e^{-\mathsf{s}}) - e^{-\mathsf{s}}\right] \tag{19}$$

which establishes that $b_j(\mathsf{S})$ is independent of σ in the log-odds space and, thus, that $b_j(S)$ is independent of q in the regular space. In the case of the informer, assuming $v_i^G = v_i^B = v_i$, the first order condition (19) simplifies into

$$a\left(1+e^{-\mathsf{S}}\right)\left(v_i + \frac{c}{r}\right) + \frac{c}{r}\left[b(1+e^{-\mathsf{s}}) - e^{-\mathsf{s}}\right] = 0, \tag{20}$$

which in the regular space can be expressed as equation (1) in the main text

$$\beta_1(s,S)\, v_i = \beta_2(s,S)\, c/r,$$

where

$$\begin{aligned}
\beta_1(s,S) &= -a(s,S)\left(1 + \frac{1-S}{S}\right) > 0 \\
\beta_2(s,S) &= b(s,S)(1 + \frac{1-s}{s}) - \frac{1-s}{s} + a(s,S)\left(1 + \frac{1-S}{S}\right) > 0.
\end{aligned}$$

Second Order Condition for the Lower Best Reply

Taking a derivative with respect to s of (18) gives

$$\begin{aligned}
&\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{s}^2} \\
= \;& \frac{e^{\sigma}}{1+e^{\sigma}} \frac{\partial \psi(\sigma,G)}{\partial \mathsf{s}} \left\{ a\left[v_j^G + e^{-\mathsf{S}} v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right] + \frac{c_j}{r}\left(b(1+e^{-\mathsf{s}}) - e^{-\mathsf{s}}\right)\right\} \\
& + \frac{e^{\sigma}\psi(\sigma,G)}{1+e^{\sigma}} \left\{ \frac{\partial a}{\partial \mathsf{s}}\left[v_j^G + e^{-\mathsf{S}} v_j^B + \left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right] + \frac{c_j}{r}\frac{\partial b}{\partial \mathsf{s}}(1+e^{-\mathsf{s}}) + \frac{c_j}{r}(1-b)e^{-\mathsf{s}}\right\}.
\end{aligned}$$

For values of s that satisfy the first order condition (19), we have

$$\left.\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{s}^2}\right|_{\mathsf{s}=b_j(\mathsf{S})} = \frac{e^{\sigma}\psi(\sigma,G)}{1+e^{\sigma}}\frac{c_j}{r}\left\{-\frac{\partial a}{\partial \mathsf{s}}\frac{1}{a}\left[b(1+e^{-\mathsf{s}}) - e^{-\mathsf{s}}\right] + \frac{\partial b}{\partial \mathsf{s}}(1+e^{-\mathsf{s}}) + (1-b)e^{-\mathsf{s}}\right\}.$$

Using

$$\begin{aligned}
1-b &= \frac{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})} - R_2 e^{R_2(\mathsf{S}-\mathsf{s})} + R_1 e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})}} \\
&= \frac{R_1 e^{R_2(\mathsf{S}-\mathsf{s})} - R_2 e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})}} = \frac{R_1 e^{-R_1(\mathsf{S}-\mathsf{s})} - R_2 e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}} = -\frac{\partial a}{\partial \mathsf{s}}\frac{1}{a},
\end{aligned}$$

the above expression simplifies to

$$\left.\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{s}^2}\right|_{\mathsf{s}=b_j(\mathsf{S})} = \frac{e^\sigma}{1+e^\sigma}\psi(\sigma,G)(1+e^{-\mathsf{s}})\frac{c_j}{r}\left[b(1-b)+\frac{\partial b}{\partial \mathsf{s}}\right],$$

which is negative, if and only if $\frac{\partial b}{\partial \mathsf{s}} < -b(1-b)$. The latter inequality can be written as

$$\frac{(R_2-R_1)^2 e^{(\mathsf{S}-\mathsf{s})}}{\left(e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}\right)^2} < \frac{\left(R_2^2+R_1^2\right)e^{(\mathsf{S}-\mathsf{s})}-R_1R_2\left(e^{2R_1(\mathsf{S}-\mathsf{s})}+e^{2R_2(\mathsf{S}-\mathsf{s})}\right)}{\left(e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_2(\mathsf{S}-\mathsf{s})}\right)^2}$$

which holds if and only if $2e^{(\mathsf{S}-\mathsf{s})} < e^{2R_1(\mathsf{S}-\mathsf{s})}+e^{2R_2(\mathsf{S}-\mathsf{s})}$. The function $e^{(1-x)(\mathsf{S}-\mathsf{s})}+e^{(1+x)(\mathsf{S}-\mathsf{s})}$ is increasing in x, so this inequality is satisfied given the expressions for R_1 and R_2. Thus, we conclude that

$$\left.\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{s}^2}\right|_{\mathsf{s}=b_j(\mathsf{S})} < 0. \tag{21}$$

(ii) Turn to the slope of the lower best reply. First, we show that $b_j(\mathsf{S}) < \mathsf{S}$ if $\mathsf{S} > \hat{\sigma}_j$ and $b_j(\mathsf{S}) = \mathsf{S}$ otherwise. We start with computing the limit of $\frac{\partial u_j(\sigma)}{\partial \mathsf{s}}$ as $\mathsf{s} \to \mathsf{S}$. Recall that

$$\frac{\partial u_j(\sigma)}{\partial \mathsf{s}} = \frac{e^\sigma}{1+e^\sigma}\psi(\sigma,G)\left\{a\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]+\frac{c_j}{r}\left[b(1+e^{-\mathsf{s}})-e^{-\mathsf{s}}\right]\right\}$$

and focus on the last term of the product. A simple calculation gives

$$\begin{aligned}&\lim_{\mathsf{s}\to\mathsf{S}}\left\{a\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]+\frac{c_j}{r}\left[b(1+e^{-\mathsf{s}})-e^{-\mathsf{s}}\right]\right\}\\ =\;&\lim_{\mathsf{s}\to\mathsf{S}}a\cdot\left[v_j^G+e^{-\mathsf{S}}v_j^B\right]-e^{-\mathsf{S}}\frac{c_j}{r}+\lim_{\mathsf{s}\to\mathsf{S}}(a+b)\cdot(1+e^{-\mathsf{S}})\frac{c_j}{r}.\end{aligned}$$

Because $\lim_{\mathsf{s}\to\mathsf{S}} a = -\infty$ and $\lim_{\mathsf{s}\to\mathsf{S}}(a+b) = 0$, one sees that the sign of the limit above depends on the sign of $v_j^G+e^{-\mathsf{S}}v_j^B$. Specifically, we have

$$\lim_{\mathsf{s}\to\mathsf{S}}\left\{a\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]+\frac{c_j}{r}\left[b(1+e^{-\mathsf{s}})-e^{-\mathsf{s}}\right]\right\}=-\infty$$

if $\mathsf{S} > \hat{\sigma}_j$, in which case $v_j^G+e^{-\mathsf{S}}v_j^B > 0$, and

$$\lim_{\mathsf{s}\to\mathsf{S}}\left\{a\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]+\frac{c_j}{r}\left[b(1+e^{-\mathsf{s}})-e^{-\mathsf{s}}\right]\right\}=+\infty$$

otherwise. Since $\lim_{\mathsf{s}\to\mathsf{S}}\frac{e^\sigma}{1+e^\sigma}\psi(\sigma,G)=\infty$, overall we have $\lim_{\mathsf{s}\to\mathsf{S}}\frac{\partial u_j(\sigma)}{\partial \mathsf{s}}=-\infty$ if $\mathsf{S} > \hat{\sigma}_j$ and $\lim_{\mathsf{s}\to\mathsf{S}}\frac{\partial u_j(\sigma)}{\partial \mathsf{s}}=\infty$ if $\mathsf{S} \le \hat{\sigma}_j$.

Next,

$$\begin{aligned}
&\lim_{\mathsf{s}\to-\infty} \frac{\partial u_j(\sigma)}{\partial \mathsf{s}} \\
= \quad &\lim_{\mathsf{s}\to-\infty} \frac{e^{\sigma}}{1+e^{\sigma}} \psi(\sigma,G) \left\{ a\left[v_j^G + e^{-\mathsf{S}} v_j^B + \left(1+e^{-\mathsf{S}}\right) \frac{c_j}{r} \right] + \frac{c_j}{r} \left[b(1+e^{-\mathsf{s}}) - e^{-\mathsf{s}} \right] \right\}
\end{aligned}$$

Focusing on the second factor, we obtain

$$\begin{aligned}
&\lim_{\mathsf{s}\to-\infty} \left\{ a\left[v_j^G + e^{-\mathsf{S}} v_j^B + \left(1+e^{-\mathsf{S}}\right) \frac{c_j}{r} \right] + \frac{c_j}{r} \left[b(1+e^{-\mathsf{s}}) - e^{-\mathsf{s}} \right] \right\} \\
= \quad &\lim_{\mathsf{s}\to-\infty} a \cdot \left[v_j^G + e^{-\mathsf{S}} v_j^B + \left(1+e^{-\mathsf{S}}\right) \frac{c_j}{r} \right] + \lim_{\mathsf{s}\to-\infty} b \cdot \frac{c_j}{r} + \lim_{\mathsf{s}\to-\infty} (b-1)\, e^{-\mathsf{s}} \frac{c_j}{r}
\end{aligned}$$

Since $\lim_{\mathsf{s}\to-\infty} \frac{e^{\sigma}}{1+e^{\sigma}} \psi(\sigma,G) = 0$, $\lim_{\mathsf{s}\to-\infty} b = R_2 > 0$ and $\lim_{\mathsf{s}\to-\infty} a = 0$, overall we have $\lim_{\mathsf{s}\to-\infty} \frac{\partial u_j(\sigma)}{\partial \mathsf{s}} = 0^+$.

Having computed the limits at the two extremes of the domain of s, we now consider two different cases. First, assume $\mathsf{S} > \hat{\sigma}_e$. Then, since $\lim_{\mathsf{s}\to\mathsf{S}} \frac{\partial u_j(\sigma)}{\partial \mathsf{s}} = -\infty$ and $\lim_{\mathsf{s}\to-\infty} \frac{\partial u_j(\sigma)}{\partial \mathsf{s}} = 0^+$, by continuity there must exist a solution to $\frac{\partial u_j(\sigma)}{\partial \mathsf{s}} = 0$, implying that in this case $b_j(\mathsf{S}) > \mathsf{S}$. Next, suppose $\mathsf{S} \leq \hat{\sigma}_j$. In this case we show that $\frac{\partial u_j(\sigma)}{\partial \mathsf{s}} < 0$. To see this assume by contradiction that there exists S_1 such that $\left.\frac{\partial u_j(\sigma)}{\partial \mathsf{s}}\right|_{\mathsf{s}=b_j(\mathsf{S}_1)} \geq 0$. Two cases can then be distinguished. If $\left.\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{s}^2}\right|_{\mathsf{s}=b_j(\mathsf{S}_1)} \geq 0$, we have a contradiction since we have shown $\left.\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{s}^2}\right|_{\mathsf{s}=b_j(\mathsf{S})} < 0$ for any S. Otherwise, since $\lim_{\mathsf{s}\to\mathsf{S}} \frac{\partial u_j(\sigma)}{\partial \mathsf{s}} = \infty$ and $\lim_{\mathsf{s}\to-\infty} \frac{\partial u_j(\sigma)}{\partial \mathsf{s}} = 0^+$, there must exist S_2 such that $\left.\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{s}^2}\right|_{\mathsf{s}=b_j(\mathsf{S}_2)} \geq 0$, which again yields a contradiction. This establishes that $b_j(\mathsf{S}) < \mathsf{S}$ if $\mathsf{S} > \hat{\sigma}_j$ and $B_j(\mathsf{S}) = \mathsf{S}$ otherwise.

Proof of Proposition 0

The Wald solution is characterized by the interior intersection of $B_w(s)$ and $b_w(S)$, which always exists by the properties established in Lemmas B1 and B2.

C Supplementary Appendix C: Additional Technical Results

Lemma C1 *The evaluator's marginal value of anticipating rejection increases in the initial belief,*

$$\frac{\partial^2 u_e}{\partial \mathsf{s} \partial \sigma_0} > 0. \tag{22}$$

Proof of Lemma C1

Using equation (18) in Appendix B for $c_j = 0$ we have

$$\frac{\partial u_e(\sigma_0)}{\partial \mathsf{s}} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \psi(\sigma_0, G) a \left[v_e^G + e^{-\mathsf{S}} v_e^B \right],$$

so that, since a does not depend on σ_0,

$$\frac{\partial^2 u_e}{\partial \mathsf{s} \partial \sigma_0} = \frac{\partial \left(\frac{e^{\sigma_0}}{1+e^{\sigma_0}} \psi(\sigma_0, G) \right)}{\partial \sigma_0} a \left[v_e^G + e^{-\mathsf{S}} v_e^B \right]. \tag{23}$$

Furthermore

$$\frac{\partial \left(\frac{e^{\sigma}}{1+e^{\sigma}} \psi(\sigma, G) \right)}{\partial \sigma} = \frac{e^{\sigma} \psi(\sigma, G) + (1+e^{\sigma}) e^{\sigma} \psi_{\sigma}(\sigma, G)}{(1+e^{\sigma})^2}$$

and

$$\psi_{\sigma}(\sigma, G) = \frac{-R_2 e^{R_2(S-\sigma)} + R_1 e^{R_1(S-\sigma)}}{e^{R_2(S-s)} - e^{R_1(S-s)}} < 0.$$

From

$$-\psi_{\sigma}(\sigma, G) = \frac{R_2 e^{R_2(S-\sigma)} - R_1 e^{R_1(S-\sigma)}}{e^{R_2(S-s)} - e^{R_1(S-s)}} > \frac{e^{R_2(S-\sigma)} - e^{R_1(S-\sigma)}}{e^{R_2(S-s)} - e^{R_1(S-s)}} = \psi(\sigma, G)$$

we have

$$\frac{\partial \left(\frac{e^{\sigma}}{1+e^{\sigma}} \psi(\sigma, G) \right)}{\partial \sigma} = \frac{e^{\sigma} \psi(\sigma, G) + (1+e^{\sigma}) e^{\sigma} \psi_{\sigma}(\sigma, H)}{(1+e^{\sigma})^2} < 0.$$

Overall, replacing in equation (23), and using $a < 0$, we obtain (22).

Lemma C2 *The evaluator's marginal value of delaying approval increases in the initial belief,*

$$\left. \frac{\partial^2 u_e}{\partial \mathsf{S} \partial \sigma_0} \right|_{\mathsf{s} = b_i(\mathsf{S})} > 0. \tag{24}$$

Proof of Lemma C2

We now consider the term $\frac{\partial^2 u_e}{\partial \mathsf{S} \partial \sigma_0}$. We proved in Appendix B that

$$\frac{\partial u_e}{\partial \mathsf{S}} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \psi(\sigma_0, G) \left[f \left(v_e^G + e^{-\mathsf{S}} v_e^B \right) - e^{-\mathsf{S}} v_e^B \right],$$

so that, since f is independent of σ_0,

$$\frac{\partial^2 u_e}{\partial \mathsf{S} \partial \sigma_0} = \frac{\partial}{\partial \sigma_0}\left(\frac{e^{\sigma_0}}{1+e^{\sigma_0}}\Psi(\sigma_0, G)\right)\left[f\left(v_e^G + e^{-\mathsf{S}} v_e^B\right) - e^{-\mathsf{S}} v_e^B\right].$$

We have

$$\frac{\partial}{\partial \sigma}\left(\frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma, G)\right) = \frac{e^{\sigma}\Psi(\sigma, G) + (1+e^{\sigma})e^{\sigma}\Psi_{\sigma}(\sigma, H)}{(1+e^{\sigma})^2} > 0.$$

Furthermore, for $S < S^N$ we have $\frac{\partial u_e}{\partial \mathsf{S}}(b_i(S), S) > 0$, so that

$$f\left(v_e^G + e^{-\mathsf{S}} v_e^B\right) - e^{-\mathsf{S}} v_e^B > 0.$$

Overall we obtain (24).

Lemma C3 *The evaluator's marginal value of delaying approval decreases in the approval standard,*

$$\left.\frac{\partial^2 u_e}{\partial S^2}\right|_{s=b_i(S)} < 0 \text{ for } S \leq S^N. \tag{25}$$

Proof of Lemma C3

From

$$\left.\frac{\partial u_e}{\partial \mathsf{S}}\right|_{\mathsf{s}=b_i(\mathsf{S})} = \frac{\partial u_e}{\partial \mathsf{s}}\frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}} + \frac{\partial u_e}{\partial \mathsf{S}}$$

we have

$$\left.\frac{\partial^2 u_e}{\partial \mathsf{S}^2}\right|_{\mathsf{s}=b_i(\mathsf{S})} = \frac{\partial^2 u_e}{\partial \mathsf{s}^2}\left(\frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}}\right)^2 + \frac{\partial u_e}{\partial \mathsf{s}}\frac{\partial^2 b_i(\mathsf{S})}{\partial \mathsf{S}^2} + 2\frac{\partial^2 u_e}{\partial \mathsf{S}\partial \mathsf{s}}\frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}} + \frac{\partial^2 u_e}{\partial \mathsf{S}^2}. \tag{26}$$

Using the expression for the utility of the evaluator

$$u_e = \frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma, G)\left(v_e^G + e^{-\mathsf{S}} v_e^B\right),$$

we now show that the four terms in the expression (26) are negative so that we have (25).

- Term 1: $\frac{\partial^2 u_e}{\partial \mathsf{s}^2}\left(\frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}}\right)^2 < 0$. From

  $$\frac{\partial^2 u_e}{\partial \mathsf{s}^2} = \frac{e^{\sigma}}{1+e^{\sigma}}\psi(\sigma, G)\left(\frac{\partial a}{\partial \mathsf{s}} + ab\right)\left(v_e^G + e^{-\mathsf{S}} v_e^B\right) < 0$$

  Simple computations yield

  $$\frac{\partial a}{\partial \mathsf{s}} + ab = a\frac{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}},$$

  from which the claim follows.

- Term 2: $\frac{\partial u_e}{\partial \mathsf{s}} \frac{\partial^2 b_i(\mathsf{S})}{\partial \mathsf{S}^2} < 0$. The utility of the evaluator is decreasing in s since the evaluator does not pay for research. The claim then follows from $\frac{\partial^2 b_i(\mathsf{S})}{\partial \mathsf{S}^2} > 0$.

- Term 3: $2\frac{\partial^2 u_e}{\partial \mathsf{S} \partial \mathsf{s}} \frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}} < 0$. Using the fact that $f\left(v_e^G + e^{-\mathsf{S}} v_e^B\right) - e^{-\mathsf{S}} v_e^B > 0$ for $\mathsf{S} < \mathsf{S}^N$, we have

$$\frac{\partial^2 u_e}{\partial \mathsf{S} \partial \mathsf{s}} = \frac{e^{\sigma}}{1+e^{\sigma}} \psi(\sigma, G) \left(fa \left(v_e^G + e^{-\mathsf{S}} v_e^B\right) - ae^{-\mathsf{S}} v_e^B \right) < 0.$$

  Given that $b_i(\mathsf{S})$ is increasing in S, the claim follows.

- Term 4: $\frac{\partial^2 u_e}{\partial \mathsf{S}^2} < 0$. From derivations above, we have

$$\frac{\partial u_e}{\partial \mathsf{S}} = \frac{e^{\sigma}}{1+e^{\sigma}} \Psi(\sigma, G) \left(f \left(v_e^G + e^{-\mathsf{S}} v_e^B\right) - e^{-\mathsf{S}} v_e^B \right),$$

  so that

$$\begin{aligned} \frac{\partial^2 u_e}{\partial \mathsf{S}^2} &= \frac{e^{\sigma}\Psi(\sigma, G)}{1+e^{\sigma}} \left[\left(f^2 + \frac{\partial f}{\partial \mathsf{S}} \right) \left(v_e^G + e^{-\mathsf{S}} v_e^B\right) + (-2f+1) e^{-\mathsf{S}} v_e^B \right] \\ &= \frac{e^{\sigma}\Psi(\sigma, G)}{1+e^{\sigma}} \left\{ f \left[f \left(v_e^G + e^{-\mathsf{S}} v_e^B\right) - e^{-\mathsf{S}} v_e^B \right] + \frac{\partial f}{\partial \mathsf{S}} \left(v_e^G + e^{-\mathsf{S}} v_e^B\right) + (1-f) e^{-\mathsf{S}} v_e^B \right\}. \end{aligned}$$

  Using the fact that $f\left(v_e^G + e^{-\mathsf{S}} v_e^B\right) - e^{-\mathsf{S}} v_e^B > 0$ for $\mathsf{S} < \mathsf{S}^N$ and that $f < 0$, we conclude $f\left(f\left(v_e^G + e^{-\mathsf{S}} v_e^B\right) - e^{-\mathsf{S}} v_e^B\right) < 0$. Given that $\frac{\partial f}{\partial \mathsf{S}} < 0$ and $1 - f > 0$ as shown above, (25) follows.