Research and the Approval Process: The Organization of Persuasion*

Emeric Henry[†] Marco Ottaviani[‡]

August 2018

Abstract

An informer sequentially collects and disseminates information through costly research to persuade an evaluator to approve an activity. Payoffs and control rights are split between informer and evaluator depending on the organizational rules governing the approval process. The welfare benchmark corresponds to Wald's classic solution for a statistician with payoff equal to the sum of informer and evaluator. Organizations with different commitment power of informer and evaluator are compared from a positive and normative perspective. Granting authority to the informer is socially optimal when information acquisition is sufficiently costly. The analysis is applied to the regulatory process for drug approval.

Keywords: Persuasion, information, organization, approval.

JEL Classification: D18 (Consumer Protection), D83 (Search; Learning; Information and Knowledge; Communication; Belief), M38 (Government Policy and Regulation).

*The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under ERC Advanced Grant 295835. We thank Alessandro Bonatti, Daniel Carpenter, Miaomiao Dong, Matthew Ellman, Marina Halac, and Joe Harrington for helpful discussion, Luca Ferocino, Marco Loseto, and Paolo Varraso for excellent research assistance, and seminar audiences at Amsterdam, Barcelona, Bergamo, Bordeaux, CEPR 2014 IMO in Frankfurt, Columbia, Copenhagen, CSEF-IGIER 2013 in Capri, EARIE 2013 in Evora, ESEM 2015 in Mannheim, ESSET 2014 and 2015 in Gerzensee, Glasgow, Innsbruck, Lausanne, MIT, Oslo, Paris, Petralia 2014, SITE 2017, Warwick, and Zurich. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Sciences Po and CEPR.

[‡]Bocconi and CEPR.

To persuade doctors or the regulatory authority to approve new drugs, pharmaceutical companies perform costly clinical trials that document safety and effectiveness. Similarly, developers can test new technologies 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 and launch the product. In these situations, a biased informer acquires and disseminates information through costly research to persuade an evaluator to adopt. The organization of the interaction between informer and evaluator varies across settings, depending on the commitment power of the two players.

For illustration, consider how the power of the evaluator increased as the drug approval process evolved in the US during the course of the twentieth century.[1] In the early days of drug regulation, pharmaceutical companies could essentially decide to take drugs to market at any point in time and let patients—or doctors on their behalf—decide on adoption. In an attempt to make verifiable the information about drug ingredients, the Pure Food and Drugs Act of 1906 introduced federal penalties for adulterating or misbranding medicines.[2] Following the public outcry over some 106 deaths caused by the use of Elixir Sulfanilamide, a drug containing a toxic solvent, Congress significantly strengthened the power of the Food and Drug Administration (FDA) by passing the Food, Drug, and Cosmetic Act of 1938. A formal approval process was established requiring drug sponsors to submit safety data to FDA officials for evaluation prior to marketing. The FDA was gradually granted the power to require further evidence if not satisfied with the initial data provided. In the wake of the tragic news of thousands of severe birth defects linked to the use of the experimental drug Thalidomide in Europe and Canada, the 1962 Kefauver-Harris Amendments introduced a number of procedures that further strengthened control of the FDA over investigational new drugs, paving the way to the modern-day FDA that can be seen as committing to an approval standard before research starts, for example by defining the margin of error.

This paper develops a flexible modeling 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 (1945) 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. Formulating the problem in continuous time to gain analytical traction, we model information collection as a stochastic process whose drift depends on a binary state, either good or bad, corre-

[1] For historical accounts see Marks (1997), Junod (2008), and Carpenter (2011).

[2] Prior to this legislation, the American Medical Association, by establishing the Council on Pharmacy and Chemistry in 1905, had already begun verifying the chemical content of drugs, forbidding advertising of misbranded drugs in its journal, and publishing a list of approved New and Nonofficial Remedies. See Cushny et al. (1905).

sponding 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. In each instant, the statistician optimally continues to conduct research provided that the posterior belief, incorporating all the information acquired up to that point, stays within the interval (s,S). The 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 falls below the lower rejection standard s.

While in Wald's classic framework a single statistician controls both research and final approval or rejection, our strategic deconstruction splits control of research and final decision between two players, an informer i and an evaluator e. The informer bears the cost of research and obtains a fixed benefit $v_i > 0$ from approval, regardless of the state. Research results are publicly disseminated.[3] 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. Wald's statistician w can be seen as a reconstructed social planner who maximizes the sum of the payoffs of informer and evaluator.

Our first contribution is the equilibrium characterization for the Wald persuasion games we introduce in Section 1. Section 1.1 sets the stage with the *informer-authority* game in which the informer is restricted to making a single take-it-or-leave-it approval request to the evaluator.[4] 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, the belief at which the evaluator is indifferent between rejection and approval when required to make an immediate decision. For intermediate beliefs the informer continues researching, trading off the expected cost of research against the approval benefit 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 from persuasion. The informer has a state-independent decision payoff and thus no interest in information per se, but still values information instrumentally in order to persuade the evaluator to approve.

[3] In the application to drug approval, information can be seen as verifiable following the establishment of the Council on Pharmacy and Chemistry in 1905 and the introduction of the Pure Food and Drugs Act of 1906.

[4] Informer authority captures the *laissez-faire* regime with verifiable information that was in place in the early days of drug evaluation after the American Medical Association established the Council on Pharmacy and Chemistry in 1905. Council members screened drugs on the basis of information presented by pharmaceutical suppliers. They could choose not to endorse the drug, but could not ask for more research. Patients and doctors had to decide whether to adopt based on the evidence available at market introduction.

Wald persuasion games combine persuasion with experimentation by incorporating into Wald's sequential analysis the strategic issues that arise when information collection and final decision are made by two different players. Our informer authority game can be seen as a continuous-time limit of Brocas and Carrillo's (2007) discrete-time model of dynamic persuasion, 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.[5]

This connection between Wald and KG is the stepping stone for our analysis. To the persuasion literature we contribute the characterization of the equilibrium in Wald persuasion games with different commitment structures. 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 1.2 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 wait for further research when not satisfied with current results.[6] 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

[5]By embedding Wald's model with binary state and constant information acquisition cost into a strategic persuasion setting, our paper characterizes a number of scenarios resulting in a static solution. As Morris and Strack (2017) show in a complementary contribution in the context of the classic (non-strategic) Wald problem, once the information acquisition cost is allowed to depend on the belief, any posterior-separable cost function (such as entropy) can be written as the outcome of a sequential sampling model.

[6]This regime resembles the system set in place by the 1938 Food, Drug, and Cosmetic Act that made it possible for the evaluator to request more information than what was voluntarily provided by pharmaceutical companies; see Marks (1997), Chapter 3. Section 3 compares the equilibrium outcomes depending on whether the FDA as evaluator maximizes consumer/patient welfare or acts as a social planner maximizing total welfare.

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, who values information but does not bear its direct cost, always sets a standard higher than the myopic cutoff. In turn, the rejection standard controlled by the informer is above the informer-authority solution. Intuitively, given that it is harder to reach the higher approval standard set by the evaluator, the informer is discouraged from acquiring information.

Section 1.3 turns to *evaluator commitment*, whereby the evaluator can commit at the outset by setting a history-independent standard that the evidence must surpass for approval to be granted.[7] As we show, the evaluator should then choose the preferred point on the informer's lower best reply. In the resulting Stackelberg outcome, the approval standard to which the evaluator commits is necessarily below the Nash level. Given that the evaluator free rides on the cost of research paid by the informer, the informer gives up on research too early from the perspective of the evaluator. Improving the informer's incentives to undertake research 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.

Our second contribution is Section 2's welfare comparison of informer authority, no commitment, and evaluator commitment. Given that the evaluator finds it optimal to commit to being 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 entails zero false negatives (type II errors), but socially excessive false positives (type I errors). In this frictionless benchmark, too many harmful drugs are approved when the pharmaceutical industry controls information. Evaluator commitment, instead, results in full information and thus attains the planner's first-best outcome, with no error of either type.

Frictions, however, loom large in typical regulatory environments. For instance, clinical trials are very costly given how difficult it is to recruit experimental subjects; furthermore, delaying approval of promising drugs damages patients. As we show, the above frictionless welfare ranking is largely overturned once we introduce information costs and/or discounting. Frictions tend to make informer

[7] Evaluator commitment could correspond to the drug approval procedure following the 1962 Kefauver-Harris Amendments of the 1938 Food, Drug, and Cosmetic Act when the "FDA was given the authority to set standards for every stage of drug testing from laboratory to clinic," according to Junod (2008). See FDA (2009) for current guidelines.

authority socially preferred to evaluator commitment, through two channels. First, when research is costly, all organizational forms lead to a higher level of false negatives compared to the social optimum—the informer always abandons research inefficiently too early; importantly, false negatives are less socially excessive under informer authority than in any other organizational form in which the evaluator retains more veto power and thus free rides more on the information.[8] Second, with sufficient discounting or research cost, false positives become socially insufficient—rather than excessive—because of the draconian standard imposed by the evaluator who does not internalize the informer's benefits from approval. With frictions, giving authority to the informer alleviates both problems: it improves the incentives to continue research when the news is bad, while also resulting in earlier approval following good news.

For our third contribution, Section 3 introduces a number of modifications of the basic model to fit the application to approval regulation. Once we add price transfers from the evaluator (the patient) to the informer (the pharmaceutical company), we compare the welfare performance of different institutions depending on the price level. We also show that under informer authority it is optimal for the planner to delegate play to consumers, provided that the price is not too high. This result can explain the social benefits of delegating approval to the professional association of doctors (representing the interests of patients) in the early days of drug regulation when government agencies (naturally mediating the interests of patients with those of the pharmaceutical industry) still did not have sufficient power to request further research. Finally, we allow for the possibility for the social planner to take up the role of evaluator in the game or, equivalently, for the government to ask the FDA to maximize social (rather than consumer/patient) welfare. Because the social planner takes into account the cost of research, we show that in some instances it becomes optimal for the planner to commit to a stricter standard of approval than under no commitment to discourage excessive testing of unpromising drugs, echoing growing concerns about the social costs of recruiting patients for experimental purposes.

Related Literature. Our model builds on the continuous-time specifications of Wald's (1945) decision-theoretic framework developed by Dvoretsky, Keifer, and Wolfowitz (1953), Mikhalevich (1958), and Shiryaev (1967); see Shiryaev (1978, Chapter IV, Sections 1-2) for a textbook treat-

[8] 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.

ment.[9] In the context of the same underlying learning model, Gul and Pesendorfer (2012) and Chan, Lizzeri, Suen, and Yariv (2018) consider strategic settings in which public information arrives over time to voters.[10] While in their setting information is revealed publicly to all voters, we focus on the sequential interaction between an informer who publicly disseminates costly information and an evaluator who makes the approval decision.

In the experimentation literature, Guo (2016) analyzes 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).[11] Assuming away pre-existing private information by the informer, we instead relax the full commitment assumption in a number 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. Instead, we consider intermediate commitment scenarios where the informer discloses the current belief when applying for approval.

In our welfare comparison of organizations in Section 2 we show that giving authority to the informer incentivizes information acquisition, resulting in a reduction of false negatives, at the ex post cost of more false positives. This result is consistent with Aghion and Tirole's (1997) key insight that delegation of authority to a biased agent may be useful to induce more information acquisition, even though this comes at the expense of inefficient decision-making ex post.

In a pioneering analysis of approval regulation, Carpenter (2004) focused on learning across different decisions. Carpenter and Ting (2007) analyzed how a firm can signal (private information about) quality to the regulator through the submission time. Orlov, Skrzypacz, and Zryumov (2016) and Bizzotto, Jesper, and Vigier (2016) consider dynamic persuasion games in which, in contrast with our work, 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.

Finally, in a contribution closely related to ours, McClellan (2017) characterizes the evaluator

[9] 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.

[10] In Gul and Pesendorfer's (2012) model, information is provided by the party that leads, whereas in Chan, Lizzeri, Suen, and Yariv (2018) voters decide collectively themselves when to stop acquiring public information and reach a decision.

[11] While the social experimentation literature spearheaded by Bolton and Harris (1999) focuses on incentives for multiple experimenters in bandit models, we focus on the interaction between a single experimenter and a decision maker. See also Strulovici (2010), who highlights how the loss of control of decision making (determined through voting in his model) reduces the incentives to acquire information and thus induces a status quo bias.

commitment solution once the restriction that the approval policy be history independent is lifted. He shows that the optimal history-dependent mechanism consists of an initial approval standard that is maintained as long as the belief remains high enough. When the belief falls below a certain level, to encourage additional research, the evaluator revises downward the approval standard, until the evaluator's myopic cutoff is reached. We maintain the assumption of history independence, given our focus on the characterization and welfare comparison of equilibrium outcomes that arise in a number of simple and realistic organizational forms. As we show, relaxing this assumption does not affect our main welfare result.

1 Wald Persuasion Games

Two risk-neutral 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 $v_i > 0$.[12]

At the outset $t = 0$ players share an initial belief about the state $q_0 = \Pr(\omega = G)$. If forced to make a decision at belief q, the evaluator approves if and only if $q \geq \hat{q}_e$, where at the myopic cutoff $\hat{q}_e$ the evaluator is indifferent between A and R,

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

At each instant the informer, instead of applying for approval, can conduct research whose results are publicly disseminated.[13] The arrival of new information is modeled as a Wiener process $d\Sigma$ with variance ρ^2 and state-dependent drift: positive drift μ in state G and negative drift $-\mu$ in state B. Acquiring information over a period of time dt costs the informer cdt; thus, the evaluator free rides on the information publicly revealed by the informer. Finally, both players discount future payoffs at the same rate $r > 0$.

[12] The analysis remains tractable if the informer's payoff is state dependent (for example because of liability or reputational losses in the bad state), but additional case distinctions must be introduced.

[13] Equivalently, we would obtain the same equilibrium outcomes if the information were privately observed by the informer but the informer were able to verifiably disclose the information acquired. This however requires that the evaluator observes how long the research lasted; see Henry (2009) as well as Herresthal (2017).

Wald Benchmark. As a welfare benchmark for comparison, consider the problem of a social planner playing the role of Wald's (1945) statistician who controls both research and approval decisions, and obtains the sum of the payoffs of the evaluator and the informer, including the cost 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:*
(i) *stops researching and rejects if $q \leq s_w^*$,*
(ii) *conducts research if $s_w^* < q < S_w^*$, and*
(iii) *stops researching and approves if $q \geq S_w^*$.*

See Supplementary Appendix B for a self-contained proof, based on the simultaneous solution of the two optimality (also known as smooth-pasting) conditions determining (i) the choice of the rejection standard $s = b_w(S)$ for a given approval standard S and (ii) the choice of the approval standard $S = B_w(s)$ for a given s.

Wald Deconstructed. Wald persuasion games deconstruct Wald's decision-theoretic model by (a) allocating the research cost to the informer, (b) splitting the payoff from approval between the players by assigning a state-independent positive payoff to the informer and the remaining payoff to the evaluator, and (c) assigning the authority over research and approval decisions to the players depending on the organization. We focus on three organizations, justified in Section 4 on the basis of the historical evolution of the drug approval process. In the first organization, informer authority, the informer can request approval only once and the evaluator cannot commit to an approval standard. We then relax these two restrictions. In the second organization, no commitment, the evaluator can request further information from the informer if not satisfied by the current evidence. In the third organization, evaluator commitment, the evaluator can also commit to an approval standard. In all these organizations, research is undertaken only if it is individually rational for the informer, who thus controls the rejection standard.

The derivation of the equilibria in these organizations builds on the best replies of informer and evaluator. To derive the informer-authority solution, Section 1.1 characterizes the informer's choice of the lower standard $s = b_i(S)$ as a best reply to a given upper standard S. Section 1.2 then characterizes the evaluator's choice of $S = B_e(s)$ as a best reply to s and derives the no-commitment solution, at which the informer's lower best reply crosses the evaluator's upper best reply. Finally, Section 1.3 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 $s = b_i(S)$.

1.1 Informer Authority

Suppose that 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. Equivalently, the informer is constrained to request approval only once. This restriction grants commitment power to the informer, because the informer is essentially committed not to doing further research after making a first proposal to the evaluator.

In the informer authority game, at each instant t the evaluator and the informer move sequentially according to the following rules:

1. The informer chooses between information acquisition $\mathscr{I}_i$, application for approval $\mathscr{A}_i$, or rejection $\mathscr{R}_i$;
2. If the informer chooses $\mathscr{A}_i$, then the evaluator chooses between approval $\mathscr{A}_e$ or rejection $\mathscr{R}_e$.
3. The outcome of the stage game is information acquisition I if $(\mathscr{I}_i)$, approval A if $(\mathscr{A}_i, \mathscr{A}_e)$, and rejection R if either $(\mathscr{R}_i)$ or $(\mathscr{A}_i, \mathscr{R}_e)$.

If the outcome of the stage game is either A or R, the game ends; following outcome I, the information acquired is publicly revealed and the stage game is played again in the next instant. We characterize the stationary Markov Perfect Equilibrium (MPE) of this dynamic game with state variable $q = q_t$.

We first derive the informer's optimal choice of rejection standard, $s = b_i(S)$, as a best reply to a given approval standard S. Denoting T_S as the first time the belief hits S with the convention that $T_S = +\infty$ if the belief hits s before S, define the expected discounted probability $\Psi(q) = E[e^{-rT_S}]$ as the expected discounted value of receiving a payoff of 1 when the belief hits for the first time the approval standard S, conditional on not hitting s before. Similarly, define the expected discounted probability $\psi(q) = E[e^{-rT_s}]$ as the expected discounted value of a payoff of 1 received when the belief hits for the first time standard s, conditional on not hitting S before.

The informer obtains a decision payoff of v_i if the approval standard S is reached before the rejection standard s; starting from q this happens with probability $\Psi(q)$. With complementary probability, the decision payoff from rejection is 0. During the research phase, the informer pays an instantaneous research cost c, and thus c/r in perpetuity, until the belief hits either s or S, which happens with probability $1 - \psi(q) - \Psi(q)$. Overall, the expected payoff of the informer is thus

$$u_i(q) = \Psi(q)v_i - [1 - \psi(q) - \Psi(q)]\,\frac{c}{r}.$$

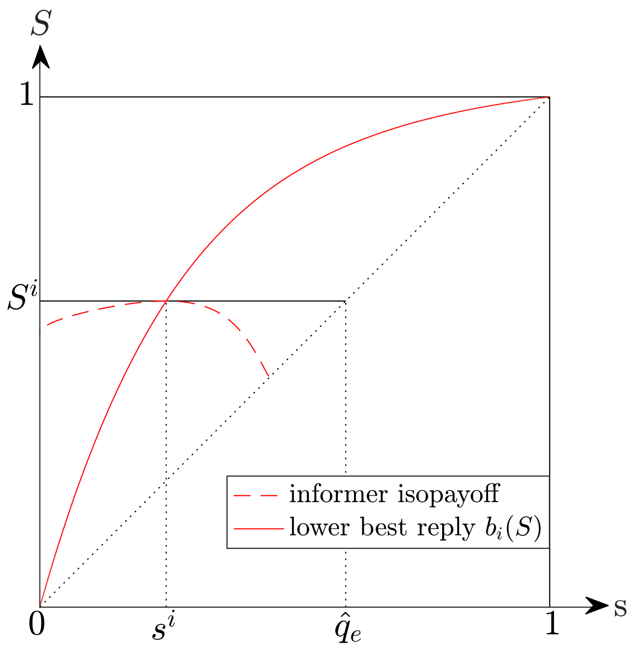

Figure 1: Informer authority solution.

At the margin, a decrease in the rejection standard increases the expected discounted probability Ψ of reaching the approval standard and obtaining v_i, but also increases the expected cost of research. The best-reply rejection standard, $s = b_i(S)$, equalizes the informer's marginal benefit of information with the marginal cost

$$\beta_1(s,S)v_i = \beta_2(s,S)\frac{c}{r}, \tag{1}$$

where $\beta_1(s,S) = -\partial\Psi/\partial s > 0$ and $\beta_2(s,S) = \partial\psi/\partial s + \partial\Psi/\partial s > 0$.[14]

The continuous red curve in Figure 1 represents the informer's lower best reply $s = b_i(S)$. To a given approval standard S, the informer best replies by choosing the rejection standard s that lies on the highest level isopayoff curve attainable at S. The figure illustrates the informer's best reply $s^i = b_i\left(S^i\right)$ against $S = S^i$: for a given prior belief q_0, the dashed red isopayoff curve is tangent to the horizontal line $S = S^i$. In the informer authority solution, $s^i = b_i\left(\hat{q}_e\right)$ given that the evaluator approves at the myopic cutoff $S^i = \hat{q}_e$:

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*

[14] On the one hand, ψ increases with s since the rejection standard s is hit earlier by the belief process. On the other hand, Ψ decreases with s since there is an increased probability that s is hit before the approval standard S. The first effect clearly dominates because an increase in s decreases the time at which the belief hits either S or s.

(s^i, S^i) such that $S^i = \hat{q}_e$ and $s^i = b_i(\hat{q}_e)$.

***(c)** Under informer authority 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 discount factor r and cost of research c converge to zero.*

For part (a.i), Figure 11 in Appendix A illustrates the construction for different values of q_0: the isopayoff curves move depending on the initial belief q_0 but, for a given approval S, the locus of tangency points remains exactly the same. For part (a.ii), Figure 12 in Appendix A illustrates the comparative statics with respect to S. 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 at a higher s. 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.[15]

According to part (b), approval is granted as soon as the belief reaches the myopic cutoff $\hat{q}_e$. The informer, who does not value information per se but only instrumentally to persuade the evaluator, stops researching as soon as there is sufficient evidence to induce the evaluator to approve. 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 shown in Figure 1.

Part (c) compares this outcome with KG:

- When the informer can commit to any signal structure, KG 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.[16] 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 the informer to choose a signal in a particular class of Brownian diffu-

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

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

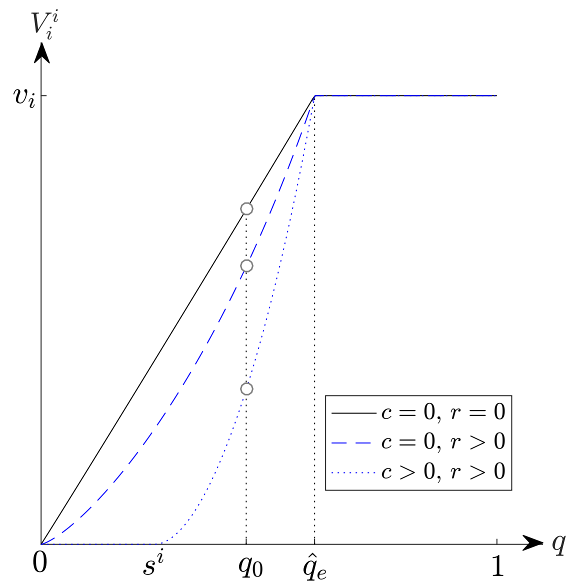

Figure 2: Informer's expected value under informer authority as a function of the belief.

> sion signals, without commitment. Nevertheless, we find that this restriction is inessential.[17] 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 at rate $r > 0$ convexifies the concavified portion of the value function where learning takes place, thus reducing the informer's payoff, as illustrated 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

[17] This observation is consistent with the following result by Morris and Strack (2017) for a single-agent version of Wald with a similar continuous-time framework: any distribution over posteriors consistent with the prior can be achieved by some experimentation strategy.

interact in a number of other realistic ways.

1.2 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 beyond the informer-authority case. As noted above, informer authority restricts the informer to submit for approval only once. Under no commitment, instead, the informer is not able to commit to not carrying out further research following rejection by the evaluator. Equivalently, in an alternative formulation a more powerful evaluator can refrain from making a decision if not satisfied with the evidence presented and can wait for more information.

Specifically, we study the stationary MPE of a dynamic game with a stage game in which to the first two substages of the informer authority game we add a third step where the informer has the option of carrying out research following rejection by the evaluator.[18] At each instant t the evaluator and the informer move sequentially according to the following rules:

1. The informer chooses between information acquisition $\mathscr{I}_i$, application for approval $\mathscr{A}_i$, or rejection $\mathscr{R}_i$;
2. If the informer chooses $\mathscr{A}_i$, then the evaluator chooses between approval $\mathscr{A}_e$ or rejection $\mathscr{R}_e$;
3. If the evaluator chooses $\mathscr{R}_e$, then the informer chooses between rejection $\mathscr{R}_i$ or information acquisition $\mathscr{I}_i$;
4. The outcome of the stage game is information acquisition I if $(\mathscr{I}_i)$ or $(\mathscr{A}_i, \mathscr{R}_e, \mathscr{I}_i)$, approval A if $(\mathscr{A}_i, \mathscr{A}_e)$, and rejection R if either $(\mathscr{R}_i)$ or $(\mathscr{A}_i, \mathscr{R}_e, \mathscr{R}_i)$.

If the outcome of the stage game is either A or R, the game ends;following outcome I, the information acquired is publicly revealed and the stage game is played again in the next instant.

The characterization of the stationary MPE of this no-commitment game with state variable $q = q_t$ relies on the construction 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 cares about the information revealed in the research process. Second, the evaluator does not bear the cost of research; nevertheless, discounting induces an opportunity cost of delaying a good decision.

[18] Alternatively, the same no commitment outcome also results in the following game: 1. The evaluator chooses between approval $\mathscr{A}_e$ or waiting $\mathscr{W}_e$; 2. If the evaluator chooses $\mathscr{W}_e$, then the informer chooses between rejection $\mathscr{R}_i$ or information acquisition $\mathscr{I}_i$; 3. The outcome of the stage game is approval A if $(\mathscr{A}_e)$, rejection R if $(\mathscr{W}_e, \mathscr{R}_i)$, and information acquisition I if $(\mathscr{W}_e, \mathscr{I}_i)$.

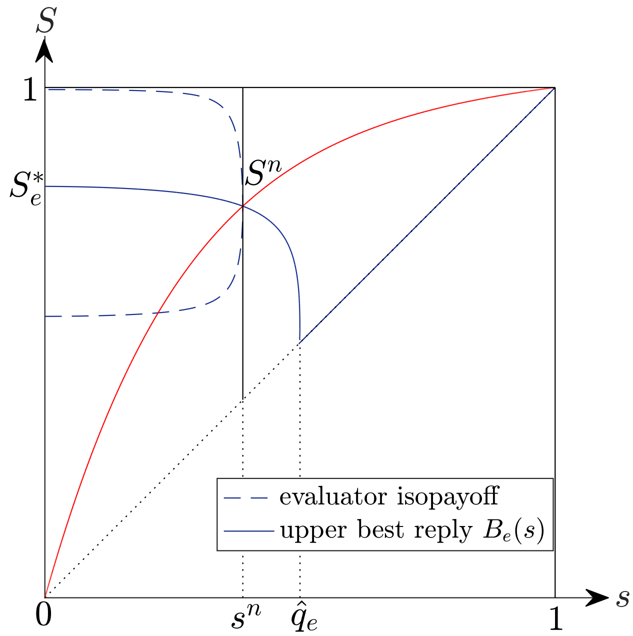

Figure 3: No-commitment solution.

For a given lower threshold s, the evaluator optimally chooses S trading off the value of information and the cost of delaying the decision. The evaluator does not take into account the cost of research and thus obtains expected payoff

$$u_e(q) = \Psi(q)\left[Sv_e^G + (1-S)v_e^B\right],$$

given that rejection results in a payoff of zero for the evaluator while approval at S results in a positive expected payoff provided that $S > \hat{q}_e$. For $s < \hat{q}_e$, the first-order condition characterizing the evaluator's best reply $S = B_e(s)$ reflects this tradeoff

$$\beta_3(s,S)\left(v_e^G - v_e^B\right) = \beta_4(s,S)\left[Sv_e^G + (1-S)v_e^B\right] \tag{2}$$

where $\beta_3(s,S) = \Psi > 0$ and $\beta_4(s,S) = -\partial\Psi/\partial S > 0$. The left-hand side represents the marginal value of information for increasing the approval standard S, while the right-hand side captures the cost of increasing S. An increase in the approval standard reduces the expected discounted probability of reaching S because of discounting and is thus costly for the evaluator.

Figure 3 shows $B_e(s)$ as the continuous blue curve, which is downward sloping for s below the myopic cutoff $\hat{q}_e$. 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. For any given rejection standard s, the evaluator chooses S on the highest isopayoff attainable at s. The figure illustrates the best reply to $s = s^n$, corresponding to a value of S such that the evaluator's isopayoff (blue dashed) is tangent to the vertical line $s = s^n$.

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 since delaying is costly, thus $B_e(s) = s$ along the diagonal.[19] 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 at a lower S. Acquisition of *less* information at the bottom induces the evaluator to require *less* information at the top—information provision is strategic complement for the evaluator, rather than substitute as for the informer.

For part (b), the unique stationary MPE outcome of this game (s^n, S^n) is displayed in Figure 3 at the intersection of the evaluator's upper best reply with the informer's lower best reply. The resulting fixed point, $s^n = b_i(B_e(s^n))$ and $S^n = B_e(b_i(S^n))$, corresponds to same outcome that would result from 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 $s^n > s^i$.[20]

[19]To see why immediate approval is optimal when $s \geq \hat{q}_e$, consider a belief $q \geq s$. By setting $S > q$, the evaluator would delay approval if S is reached and run the risk of rejection if instead s is reached, at a belief where the evaluator would in fact prefer to approve.

[20]Given that the approval standard is set at a level further away from the unconstrained optimal level for the informer, $S_i^* = 0$, the value of research at the bottom for the informer is reduced by Proposition 1.a.ii.

1.3 Evaluator Commitment

The evaluator is weakly better off under no commitment than under informer authority, and strictly better off for a prior in (s^n, S^n).[21] The evaluator can do even better by initially committing to approve according to an ex-ante specified rule. 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 as described below:[22]

Proposition 3 *In the evaluator-commitment game:* ***(i)*** *if $q_0 \in (s^i, S^n)$ the evaluator chooses an interior commitment $S^e(q_0) \in (S^i = \hat{q}_e, S^n)$ increasing in q_0;* ***(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.*

According to part (i), the evaluator encourages the informer to perform more research at the bottom by committing to an approval standard below the no-commitment level S^n, trading off:

- *A second-order negative direct effect:* For a given rejection standard s, the evaluator suffers a loss when reducing the approval standard marginally below $S = B_e(s)$. This loss is second order by the envelope theorem.
- *A first-order positive strategic effect:* By reducing S, the evaluator induces the informer to decrease s, since $b_i(S)$ is increasing. Given that s is chosen by the informer at a level that is strictly higher than the evaluator's optimal $s_e^* = 0$ and given that the evaluator's payoff strictly decreases in s, the induced reduction in s results in a first-order gain for the evaluator.

To encourage research by the informer, the evaluator optimally resolves this tradeoff by committing to an approval standard S^e at the tangency point between the highest isopayoff of the evaluator and the informer's best reply $b_i(S)$. Figure 4 illustrates the tangency with the evaluator's isopayoff curve (dashed curve) for an initial belief fixed at q_0.[23] When interior, the optimal commitment is lenient,

[21] Under informer authority, the evaluator obtains no value of information, while under no commitment the evaluator enjoys valuable learning for $q_0 \in (s^n, S^n)$ and does not bear its cost.

[22] McClellan (2017) also characterizes the two-sided commitment solution by which the evaluator makes the informer just indifferent between accepting or rejecting at date 0; this solution corresponds to the evaluator choosing the preferred point on the informer's zero-level iso-payoff in Figure 1.

[23] The shape of the iso-payoff curve depends on q_0

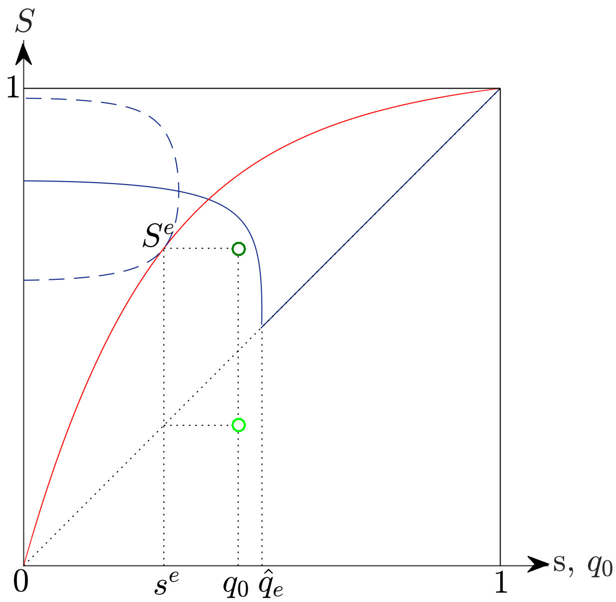

Figure 4: Commitment solution.

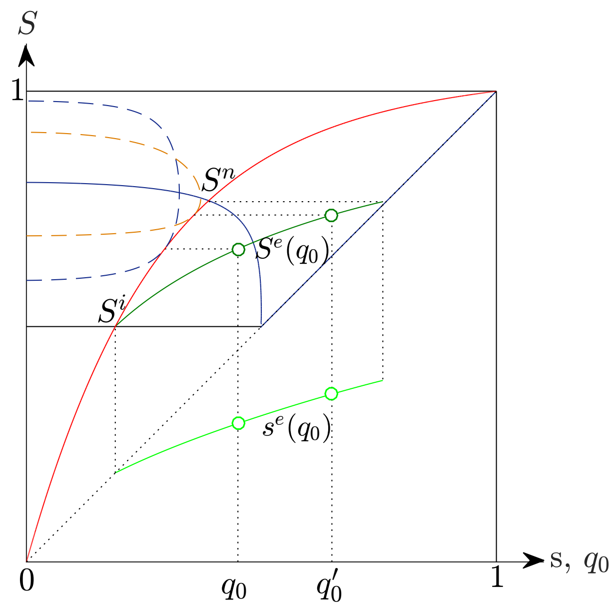

Figure 5: Commitment path.

with $S^e < S^n$. Intuitively, as in Aghion and Tirole (1997), the evaluator relinquishes some "real" authority to the informer, so as to improve the incentives to acquire valuable information.

The commitment solution depends on the initial belief $q_0 \in (s^i, S^n)$ at the time of commitment. The upper curve plotted in dark green in Figure 5 traces the commitment path $S^e(q_0)$ as a function of the initial belief q_0 (rather than s) on the horizontal axis; the lower curve plotted in light green corresponds to the resulting $s^e(q_0) = b_i(S^e(q_0))$. To understand the construction, fix an initial belief $q_0 \in (s^i, S^n)$ and find the commitment solution S^e as a function of q_0 by projecting the tangency point S^e to the vertical dotted line at q_0, as in Figure 4. The commitment path $S^e(q_0)$ is obtained by repeating this procedure for each value of q_0, as in Figure 5.

As expressed in Proposition 3.(i), an increase in the initial belief from q_0 to q_0' results in an increase in the commitment solution from $S^e(q_0)$ to $S^e(q_0')$. Intuitively, an increase in the initial belief makes the commitment to any S below S^n more costly (because the expected discounted probability of reaching a given ex-post suboptimal approval standard is increased) and reduces the strategic benefit (because the expected discounted probability of reaching the rejection standard chosen by the informer is reduced). An increase in the initial belief, $q_0' > q_0$, induces a clockwise rotation in the evaluator's isopayoff curve (in the relevant range, below the evaluator's best reply), resulting in a tangency at $S^e(q_0') > S^e(q_0)$ with the informer's best reply, as shown in Figure 5. When interior, the commitment path is strictly increasing in q_0, starting from $S^e = S^i = \hat{q}_e$ for $q_0 = s^i$ and ending at $S^e = S^n$ for $q_0 = s^n$. Thus, the commitment solution is not dynamically consistent.

The optimal commitment is above the myopic cutoff $\hat{q}_e$, because otherwise the evaluator would obtain a negative payoff, which can be attained by always rejecting. When the initial belief is below s^i, as in part (ii), any commitment above $\hat{q}_e$ will push the informer to do no research given that $b_i(S) > q_0$ for $S > \hat{q}_e$, so that in this region commitment has no value for the evaluator.[24] At the other extreme, as in part (iii), when the prior belief is above S^n, it is optimal for the evaluator to choose immediate approval, thereby inducing the informer to do no research.[25]

For simplicity, this paper restricts attention to history-independent approval mechanisms. McClellan (2017) shows that the optimal history-dependent mechanism consists of an initial approval standard that is maintained at a fixed level until the posterior reaches the informer's lower best reply to it. Whenever the lower best reply is hit, the informer, who would otherwise reject, is kept experimenting by a gradual reduction of the approval standard, until the myopic cutoff is reached. The optimal history-dependent approval mechanism does not depend on the initial belief.

2 Welfare Comparison of Organizations: Wald Reconstructed

We now turn to our second contribution, the welfare comparison of informer authority, no commitment, and evaluator commitment. The normative benchmark is Proposition 0's Wald solution for the social planner who controls both research and approval decisions with the objective of maximizing the sum of the payoffs of informer and evaluator, $v_w^\omega = v_e^\omega + v_i$, also bearing the cost of research. The Wald solution (s_w^*, S_w^*) is at the intersection between the planner's optimality conditions $S = B_w(s)$ and $s = b_w(S)$—pink and blue dashed-dotted curves in Figure 7— representing the optimal choice of a standard for a given choice of the other standard.

How do the planner's optimality conditions compare to the best replies of evaluator and informer? The planner's upper optimality condition $B_w(s)$ always lies below the evaluator's upper best reply $B_e(s)$, for two reasons. First, the planner internalizes the informer's cost of research and thus has an incentive to approve earlier than the evaluator. Second, the planner internalizes the informer's positive payoff from approval, and thus the expected value of approval for the planner is always higher than for the evaluator—indeed, the planner's myopic cutoff $\hat{q}_w$ is below the evaluator's myopic cutoff $\hat{q}_e$, and more generally the planner has a higher opportunity cost of research because of the informer's benefit from approval.

[24]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 in which the standard would be set exactly at the myopic cutoff.

[25]Thus, the set of initial beliefs q_0 for which some research is conducted under evaluator commitment is larger than under no commitment: $(s^n, S^n) \subset (s^i, S^n)$.

The upper best reply $B_w(s)$ is increasing for $s \leq s_w^*$, it reaches its maximum at the planner's stand-alone solution, s_w^*, to then decrease 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, so that it becomes optimal for the planner to increase the approval standard toward the Wald solution S_w^*; in this case, information at the bottom and at the top are substitutes. When instead the rejection standard is excessively tough ($s > s_w^*$), the planner responds to a higher rejection standard by decreasing the approval standard further below S_w^*; then, information at the bottom and at the top are complements. Finally, when the rejection standard is above the myopic cutoff, $s > \hat{q}_w$, it is optimal for the planner to immediately stop information acquisition; thus, the upper optimality condition lies along the diagonal, $B_w(s) = s$.

Turning to the lower best reply $b_w(S)$, note that the addition of the evaluator's approval payoff to the constant payoff of the informer makes the planner's payoff state-dependent. Unlike the informer, the planner values research per se because information about the underlying state of the world is valuable. For $S > \hat{q}_w$, the planner's lower optimality condition 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.[26] 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 optimality condition then lies along the diagonal, $b_w(S) = S$.

How do the lower best replies of planner and informer compare? The planner is more (respectively less) eager to research than the informer

$$b_w(S) \lesseqgtr b_i(S) \Leftrightarrow S \gtreqless \hat{q}_e$$

whenever the approval standard is above (respectively below) the evaluator's myopic cutoff. Intuitively, when approval benefits (respectively hurts) the evaluator, the value of information at the bottom for the planner is higher than for the informer. Given that the evaluator's approval standard will always be above $\hat{q}_e$, whenever information is costly ($c > 0$) all decentralized organizations result in too little information acquisition at the bottom compared to the socially optimal level.[27]

[26] 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 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^*; information at the top is then substitute for information at the bottom.

[27] This key property is lost when the informer plays with the social planner, as in Section 3.2.

2.1 False Positives and False Negatives

Equilibrium in different organizations results in a balance between false negatives (type II errors, determined by the rejection standard s) and false positives (type I errors, determined by the approval standard S):

- First, by Proposition 3 we can rank organizations in terms of false positives: given that under evaluator commitment the evaluator commits to reducing the standard but still keeps it above the myopic cutoff, informer authority leads to more false positives than evaluator commitment, which in turn leads to more false positives than no commitment, $S^i \leq S^e \leq S^n$.
- Second, this ranking for false positives naturally translates into a ranking for false negatives, $s^i \leq s^e \leq s^n$. Given that the informer's lower best reply increases in S, rejection and approval standards move in the same direction. As the approval standard becomes tougher, the informer is discouraged from undertaking costly research and abandons for more favorable beliefs. Thus, organizations that lead to more false positives also result in fewer false negatives.

But how do the errors resulting in different organizations compare to the socially optimal levels? To set the stage, consider the benchmark with costless research ($c = 0$) and no discounting ($r = 0$), corresponding to KG's setting. In this frictionless environment, the socially optimal solution entails full information, with neither false positives nor false negatives. In KG's solution (corresponding to informer authority) the informer submits for approval as soon as the evaluator's myopic cutoff is reached, thus resulting in socially excessive false positives—the informer obtains the highest acceptance probability compatible with rationality by the evaluator, but does not internalize the evaluator's value of additional information. Evaluator commitment, instead, eliminates false positives altogether. Given that without frictions all organizations result in zero false negatives, $s^*_w = s^i = s^e = s^n = 0$, evaluator commitment (as well as no commitment) dominates informer authority and achieves the social optimum, $S^i < S^*_w = S^e = S^n = 1$. We now consider how the results change when we depart from the frictionless benchmark:

Proposition 4 ***(a)*** *False negatives are socially excessive in all decentralized organizations, higher in no commitment than in evaluator commitment, and higher in evaluator commitment than in informer authority:* $s^*_w \leq s^i \leq s^e \leq s^n$*, with strict inequalities whenever* $c > 0$.

(b) *There exists* $\bar{c}$ *such that:*

(i) *if* $c \geq \bar{c}$*, false positives are socially insufficient under informer authority and even more so under evaluator commitment,* $S^*_w \leq S^i \leq S^e \leq S^n$*;*

(ii) *if $c < \overline{c}$, false positives are socially excessive under informer authority $S^i \leq S_w^*$. Furthermore, there exists an initial belief $\overline{q}_0(c)$ such that false positives are socially insufficient under evaluator commitment, $S^e \geq S_w^*$, if and only if $q_0 \geq \overline{q}_0(c)$.*

Research costs introduce two forces that tend to overturn the frictionless welfare comparison. Part (a) highlights the first force. Regardless of the specific organizational form, with costly research decentralized interaction between informer and evaluator hinders research at the bottom and leads to excessive rejection. Intuitively, the evaluator has veto power over the approval decision and therefore only approves when obtaining a positive payoff. Because research carries a positive option value that is not internalized by the informer, research stops too early and false negatives are socially excessive—the more so the more exacting the approval standard required by the evaluator.[28] Thus, informer authority unambiguously dominates evaluator commitment in terms of false negatives.

Part (b) compares the extent of false positives in different organizations relative to the socially optimal level, uncovering a second force that tends to favor informer authority over evaluator commitment. When the cost of research is sufficiently high, $c > \overline{c}$, as in part (b.i) and Figure 6, informer authority dominates evaluator commitment in terms of false positives. Intuitively, the approval standard under informer authority $S^i = \hat{q}_e$ is independent of c. In contrast, for the social planner approval depends on c. As c increases, the social planner tolerates more false positives, with S_w^* converging to $\hat{q}_w$ as $c \rightarrow +\infty$, where $\hat{q}_w$ is the planner's myopic cutoff. As a consequence, since $\hat{q}_w < \hat{q}_e$, for sufficiently high c, false positives become socially insufficient under informer authority, $S_w^* < S^i$. Informer authority then dominates evaluator commitment which results in even fewer false positives by granting more veto power to the evaluator, $S_w^* < S^i \leq S^e \leq S^n$.[29]

Turning to part (b.ii), when the cost of research is sufficiently small, $c < \overline{c}$, false positives are still socially excessive under informer authority as in the frictionless world with $c = 0$. The comparison between evaluator commitment and the social optimal level then depends on the initial belief q_0, as illustrated by Figure 7. As we know from Proposition 3.i, the commitment path is an increasing function of q_0 taking values in $\left[S^i, S^n\right]$, while the socially optimal level S_w^* is independent of q_0. When the initial belief is at the boundary level $q_0 = \overline{q}_0$, evaluator commitment exactly results in socially optimal approval, $S^e(\overline{q}_0) = S_w^*$. For initial beliefs below $\overline{q}_0$, the evaluator commits to a standard below the socially optimal level, so false positives are socially excessive $S^i \leq S^e \leq S_w^* < S^n$; evaluator commitment then performs better than informer authority also in terms of false positives.

[28]To prove this result we also need the fact that for $S > \hat{q}_e$ the minimum of the lower best reply is obtained for $s = s_w^*$.

[29]Following a similar logic to part (b), for any fixed c, there also exists $\overline{r} > 0$ such that false positives are socially excessive $S^i \leq S_w^* < S^e < S^n$ (resp. insufficient $S_w^* \leq S^i < S^e < S^n$) under informer authority if $r \leq \overline{r}$ (resp. $r \geq \overline{r}$).

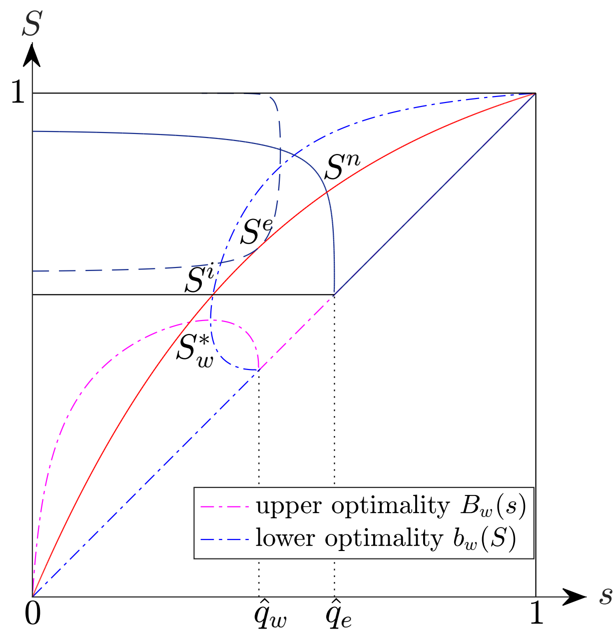

Figure 6: Comparison of solutions for $c > \bar{c}$.

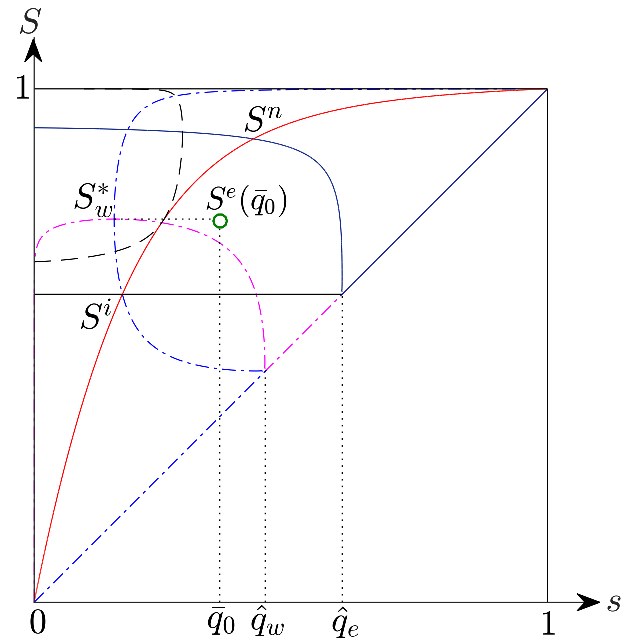

Figure 7: Comparison of solutions for $c < \bar{c}$.

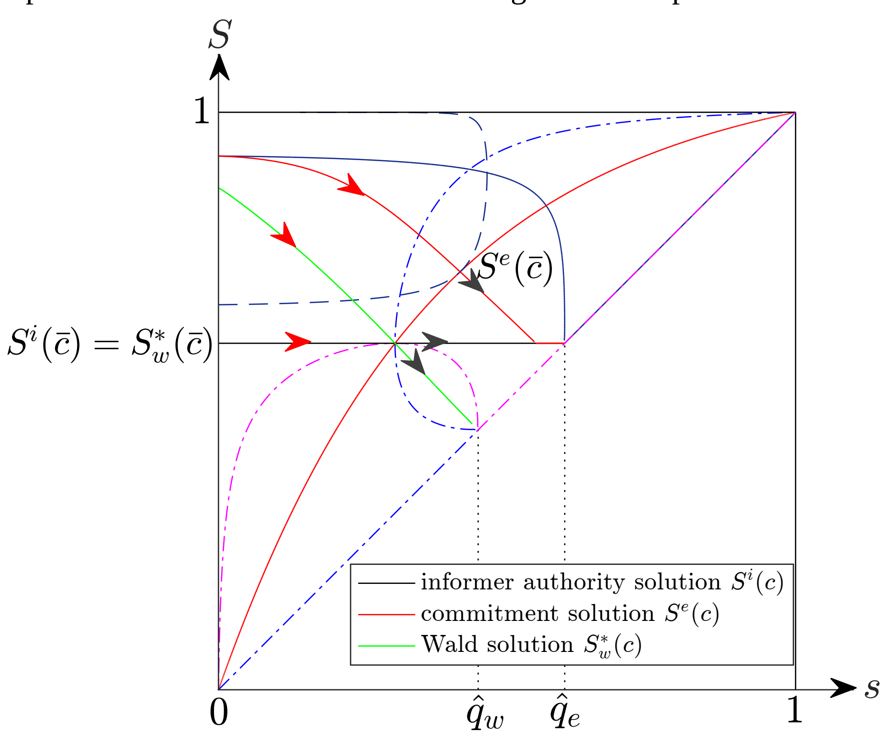

Figure 8: Comparative statics of solutions as c increases.

For initial beliefs above $\overline{q}_0$, instead, the resulting commitment is higher than S_w^*, leading to a socially insufficient level of false positives.

How small is the boundary level of research cost $\overline{c}$ at which $S^i = S_w^*$? As we show in the proof of Proposition 4, $\overline{c} = 0$ when the discount rate r is sufficiently large. In this case, false positives are socially insufficient under informer authority, and a fortiori under the other organizational forms. The evaluator's veto power leads to delays that are excessively costly in any organization.

Figure 8 illustrates the comparative statics of the solutions $S_w^*(c)$, $S^i(c)$, $S^e(c)$ as the research cost c increases; best replies and optimality conditions correspond to $c = \overline{c}$.[30] Holding q_0 fixed, the red arrows coincide with the solutions from Figure 7 for $c < \overline{c}$, whereas the black arrows correspond to the solutions in Figure 6 for $c > \overline{c}$. The informer authority solution moves to the right along the horizontal line $\left(s^i, \hat{q}_e\right)$, the social planner solution (s_w^*, S_w^*) moves downward and to the right toward the planner's myopic cutoff $(\hat{q}_w, \hat{q}_w)$, while the evaluator commitment solution (s^e, S^e) moves downward and to the right toward the evaluator's myopic cutoff $(\hat{q}_e, \hat{q}_e)$. When $c = \overline{c}$, the social optimum path $S_w^*(c)$ intersects the informer authority path $S^i(c)$; in this case, informer authority results in the socially optimal level of false positives.

2.2 Welfare Comparison

Combining these effects on false positive and false negatives, we now show that informer authority socially dominates the other organizations as long as c and r are not too low, thus overturning the frictionless comparison for a large set of parameters.

Proposition 5 ***(a)*** *Evaluator commitment Pareto dominates no commitment.*
(b) *For any r there exists $\tilde{c} < \overline{c}$ such that if $c > \tilde{c}$ informer authority welfare dominates evaluator commitment.*

Part (a) focuses on the welfare comparison between no commitment and evaluator authority. By revealed preference, the evaluator must benefit from commitment as a Stackelberg leader—the no commitment solution is always a possible choice. Given that the approval standard resulting with commitment is more lenient by Proposition 3, the informer also gains from the increased probability of acceptance. Overall, evaluator commitment Pareto dominates no commitment.

Part (b) shows that informer authority welfare dominates evaluator commitment if c is high enough. According to Proposition 4.a and 4.b.i, if $c > \overline{c}$, informer authority clearly dominates

[30] If q_0 was lower, the three paths would still be decreasing, but paths $S_w^*(c)$ and $S^e(c)$ would cross for an intermediate value of c.

evaluator commitment in terms of both false positives and false negatives. Proposition 5.b shows that this result extends to lower costs: there exists $\tilde{c} < \overline{c}$, such that if $c \in (\tilde{c}, \overline{c})$ informer authority is also socially preferred: the positive effect of limiting false negatives dominates the potential social cost of having too many false positives.

To summarize, the result obtained in the frictionless benchmark remains valid only for a restricted range of parameters, namely for c and r sufficiently low. For instance, if r is sufficiently high then $\tilde{c} = 0$, so that informer authority dominates evaluator commitment for all costs, as shown in the proof of Proposition 5. Once research is costly, false negatives come into the picture, thus favoring informer authority. Discounting implies that false positives are increasingly desirable as delays are costly. As c or r increase, the preferred organization is informer authority, the organization that grants the smallest veto power to the evaluator, who does not take into account the informer's benefit from approval and free rides on the cost of research borne by the informer.[31]

Our key insight that informer authority dominates evaluator commitment when c is relatively high remains valid even if we allow the evaluator to commit to McClellan's (2017) history-dependent approval mechanism. To illustrate this point, consider the history-independent commitment S^e for $c > \overline{c}$ as in Figure 6. With history dependence, if the belief falls below $b_i(S^e)$ following bad news, the evaluator decreases the approval standard, but never below the myopic cutoff $S^i = \hat{q}_e$, where the approval payoff becomes negative. Thus, the informer authority solution still lies closer to the social optimum and still welfare dominates.

3 Approval Regulation

To apply the model to approval regulation, this section extends the baseline analysis as follows: Section 3.1 introduces a price P transferred from the evaluator to the informer. Section 3.2 considers the case where the social planner (FDA in case of drug approval) takes the role of evaluator, either under informer authority in Section 3.2.1 or evaluator commitment (that becomes planner commitment) in Section 3.2.2.

3.1 Role of Prices

When applying the model to approval regulation for products such as drugs, the players' payoffs depend on the product price, P. For example, pharmaceutical companies conduct research and,

[31] Results similar to Proposition 5 can be derived with respect to variations in the discount rate and the informer's payoff from approval. Informer authority welfare dominates evaluator commitment if v_i or r are sufficiently large.

upon approval, collect P when selling the drug. The pharmaceutical company plays the role of informer with approval payoff $v_{i,P}^{\omega} = P$.[32] Patients obtain a state-dependent payoff v_e^{ω} and pay $P \leq v_e^G$ to purchase the drug. Patients—or doctors acting on their behalf—play the role of evaluator with approval payoff $v_{e,P}^{\omega} = v_e^{\omega} - P$.[33]

We assume that the price P is set exogenously at the start of the game. This case is relevant when modeling drug prices that are set through a regulatory negotiation process, 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).[34]

Proposition 6 *There exists a threshold price level $\tilde{P}$ such that informer authority is strictly socially preferred to evaluator commitment if and only if $P \geq \tilde{P}$.*

Proposition 6 shows that the price level determines whether informer authority or evaluator commitment is preferred. The socially optimal outcome is clearly invariant with respect to the price P, which is a simple transfer. As P increases, the informer's lower best reply shifts to the left and the evaluator's upper best reply shifts up, as the evaluator with veto power becomes more reluctant to approve. When P is large, this veto power is reduced, making informer authority preferable.[35]

What if instead players were allowed to set the price later in the game? First, if the informer endogenously sets the price when requesting approval, we argue that the first-best solution is achieved regardless of the organizational form. Following the logic of the Coase theorem, the informer is able to extract the evaluator's full surplus regardless of the organizational form. The informer's incentives become perfectly aligned with those of the social planner, so that the organization of persuasion becomes irrelevant.[36] Second, if instead the evaluator sets the price at the time when approval is requested, the evaluator expropriates the informer. Because of this hold up problem, the informer's incentives to collect information are eliminated altogether. Thus, in this second case, no

[32] Pharmaceutical companies are largely shielded from liability unless they are shown to have misrepresented evidence. The analysis can be easily extended to allow the informer to have a state-dependent payoff, capturing liability or reputation costs incurred if the drug is shown to be unsafe.

[33] Doctors could constitute an additional category of actors. However, as a first approximation we assume away conflicts of interest between doctors and patients—and thus identify patients with doctors.

[34] For example, in the US drug prices are mostly negotiated by insurance companies, while drug approval is coordinated by the FDA.

[35] There exists a value $\tilde{P}$ such that, if P is set at $\tilde{P}$ before the start of the game, the social optimum is attained under informer authority. Indeed, the informer authority outcome $(s^i(P), S^i = \hat{q}_e(P))$ lies on the planner's lower best reply curve, given that for $S = \hat{q}_e$ there is no externality, $b_w(S) = b_i(S)$. As P is raised, $\hat{q}_e(P)$ increases and the informer-authority outcome (s^i, S^i) moves up along the planner's lower best reply curve b_w. The price P can thus be chosen so that the intersection of $b_w(\hat{q}_e(P))$ and $b_i(\hat{q}_e(P))$ is exactly at the social optimum.

[36] 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).

information is collected, regardless of the organizational form. In the rest of the paper we focus on the case where P is set exogenously.

3.2 Social Planner as Active Player

In the baseline model, the approval decision is taken by an evaluator who does not internalize the payoff of the informer. In the application to the drug approval process, this assumption undoubtedly fits well the early days when doctors/patients/consumers played the role of evaluators. As the authority to approve shifts to the FDA, the objective of the agency becomes relevant. If the agency maximizes the welfare of consumers, the conflict highlighted in the analysis so far remains relevant. However, results change if the objective of the FDA is to maximize social welfare.

We now extend the analysis to the case in which the FDA acts as a social planner, taking into account the welfare of both the consumer and the pharmaceutical company. Given that $v^{\omega}_{i,P} = P$, the planner then obtains decision payoff $v^{\omega}_{w} = v^{\omega}_{e,P} + v^{\omega}_{i,P} = v^{\omega}_{e}$ upon approval, while bearing the research cost during the learning phase. Section 3.2.1 gives conditions for when the planner benefits from delegating play to the evaluator in the informer authority game. Section 3.2.2 then characterizes the planner commitment solution in which the planner commits to a history-independent approval standard.

3.2.1 Strategic Delegation of Evaluation

We now characterize when the planner prefers to commit not to be an active player in the informer authority game by strategically delegating approval to the evaluator. To this end, denote with upper (rather than lower) case superscripts the outcomes of the game between informer and planner (rather than evaluator, as in baseline); for example, $s^I(P) = b_{i,P}(\hat{q}_w)$ is the rejection standard in the informer authority game against the planner with exogenous price P.

Proposition 7 *Under informer authority, the planner strictly benefits from delegating play to the evaluator if:*
(i) *the starting belief is low,* $q_0 \in (s^I(P), \hat{q}_w)$ *or*
(ii) *the starting belief is intermediate* $q_0 \in [\hat{q}_w, S^*_w)$ *and the price is low* $P \in (\check{P}(q_0), \hat{P}(q_0))$.

This result extends a logic familiar from the delegation literature at least since Fershtman and Judd (1987).[37] In case (i), the planner strictly benefits from delegating to the evaluator when q_0

[37] See also Dessein (2002) and Armstrong and Vickers (2010) for related insights.

is below the planner's myopic cutoff $\hat{q}_w$ but above $s^I(P)$. Intuitively, for these starting beliefs the informer conducts research and submits for approval as soon as the myopic cutoff of the planner $\hat{q}_w$ is reached and as long as the posterior remains above $s^I(P)$. Note that the social payoff is zero at approval by definition of the myopic cutoff. Delegation to the tougher evaluator thus allows the planner to forestall costly but socially worthless research.

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. If P becomes too high, the planner wants to retain control, to prevent socially excessive research.

This result is consistent with the evolution of drug approval regulation in the US, where in the early days before the government would gain substantial power through the Food, Drug, and Cosmetic Act of 1938 and, especially, the 1962 Kefauver-Harris Amendments, drug regulation was largely delegated to patients and the professional association of doctors, the American Medical Association, naturally representing the interests of patients.

3.2.2 Planner Commitment

According to Proposition 7, under informer authority delegating play to the evaluator is welfare improving when the informer is poorly motivated. By revealed preference, however, delegation is unambiguously suboptimal when the player in charge of the approval decision has commitment power. Indeed, as government gained commitment power in the second half of the twentieth century, it took directly take charge of setting standards for drug approval regulation. Extending the analysis of evaluator commitment in Section 1.3, we now characterize the planner's commitment solution.

As shown in Section 1.3, the evaluator optimally commits to an approval standard below the Nash level to encourage research by the informer at low beliefs. Given that the cost of research is borne only by the informer, the evaluator always benefits from setting a lenient commitment so as to induce the informer to do more research at the bottom. The planner, instead, also cares about the cost of research. Thus, when the informer is strongly motivated, for instance because the price is high, it is optimal for the planner to commit to an approval standard above the Nash level (resulting in the no-commitment game between informer and planner), so as to discourage the informer from carrying out excessive research.

Denoting by P^* the price at which the no-commitment Nash solution is socially optimal and using capitalized superscripts for the outcomes of the games between informer and planner, the planner's optimal commitment $S^W(q_0)$ is characterized as follows:

Proposition 8 *If the price paid to the informer is relatively low, $P < \overline{P}$, 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: $S^W > b_i^{-1}(q_0)$.*
(b) *For $q_0 \in (\tilde{q}, S^N)$, the planner chooses an interior commitment such that:*

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

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

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

(c) *For $q_0 > S^N$, the planner approves immediately: $S^W(q_0) \leq q_0$.*

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), at the Nash outcome the informer undertakes socially insufficient research. 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 $\lim_{q_0 \to S^{N-}} S^W(q_0) = S^N$, with the planner commitment solution converging to the Nash level along an upward sloping path, as shown in Figure 14 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 $q_0 \to S^N$, 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 15 in Appendix A.

Case (iii) highlights a discontinuity in the commitment path when $P \in (\underline{P}, \overline{P})$. For P slightly above $\underline{P}$, the informer's lower best reply (red curve in Figure 9) crosses the planner's zero-level isopayoff 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 isopayoff (continuous black) is tangent to the informer's lower best reply at a belief, $b_i(S^W(\tilde{q}))$, 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, $S = S^W(\tilde{q})$, and a blocking commitment, corresponding to any $S \geq b_i^{-1}(\tilde{q})$. Since $b_i(S^W(\tilde{q})) < \tilde{q}$, such an interior commitment induces the informer to set the rejection standard below $\tilde{q}$. The planner's optimal interior commitment $S^W(\tilde{q})$ must be strictly below the lowest blocking commitment $b_i^{-1}(\tilde{q})$ and this gives rise to the discontinuity displayed in the figure.[38]

These results shed light on approval regulation in the pharmaceutical sector. As outlined above, with the 1962 Drug Amendments the FDA was given the authority to commit to standards of approval. In some circumstances, the FDA should choose a blocking commitment, along the lines of case (a). Indeed, before being allowed to start clinical trials on humans, pharmaceutical companies need to submit an Investigational New Drug 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 clinical testing.[39]

Cases (ii) and (iii) correspond to situations in which the FDA enacts a relatively tough commitment when allowing clinical testing. In these cases pharmaceutical companies have excessive incentives to undertake clinical trials to prove effectiveness. Tough commitment serves the purpose of discouraging pharmaceutical companies from running costly clinical trials resulting in information with high private value for the pharmaceutical companies but relatively low social value.[40]

[38]To understand the discontinuity, consider the dashed curve in Figure 9, which represents the value function of the planner for $S = S^W(\tilde{q})$, as a function of the initial belief q_0. As illustrated in the figure, at $q_0 \leq b_i(S^W(\tilde{q}))$ 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 $q_0 \in (b_i(S^W(\tilde{q})), \tilde{q})$ 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 $b_i(S^W(\tilde{q})) < \tilde{q}$, such range of beliefs is not empty so that a discontinuity arises.

[39]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.

[40]This problem can be prevalent with the arrival of new treatment strategies such as immunotherapy. The *New York*

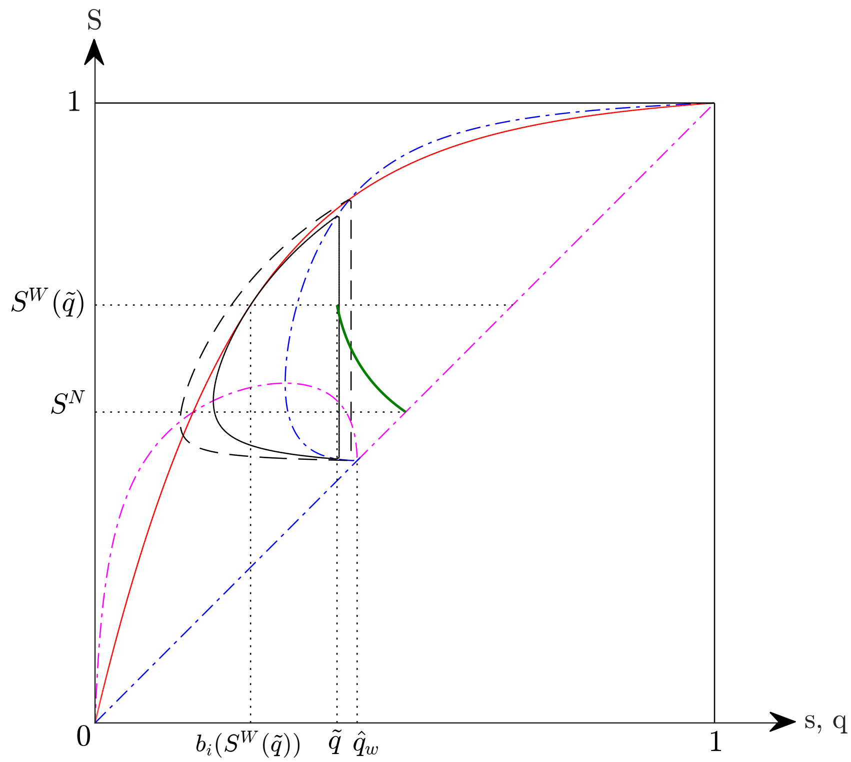

Figure 9: Discontinuous commitment path.

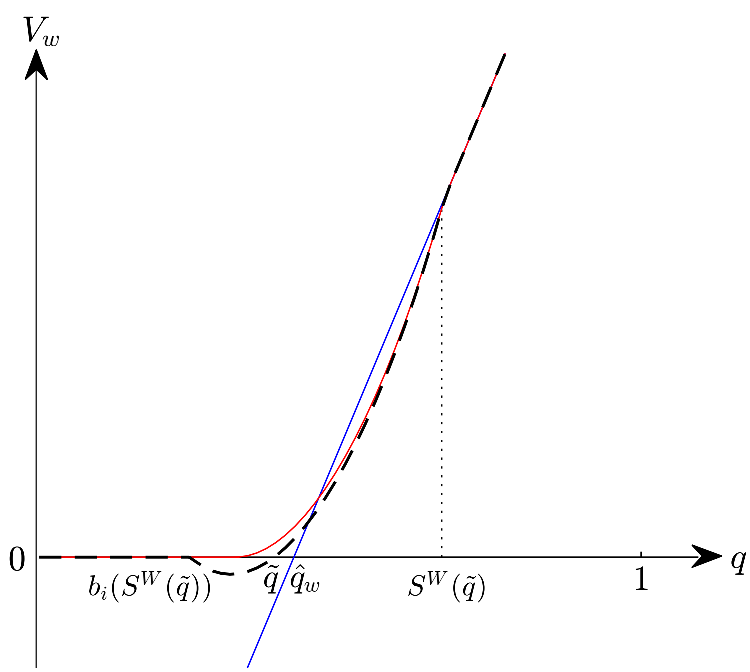

Figure 10: Planner value function.

4 Conclusion

Thanks to its analytical tractability, this continuous-time model with binary state can be extended in a number of realistic directions. Adding initial private information by the informer, McClellan (2017) shows that in the optimal mechanism the high type informer is given initially a commitment to a lower approval threshold but is penalized following release of unfavorable public information.[41] To investigate the optimal mix of ex-ante and ex-post safety regulation, Henry, Loseto, and Ottaviani (2018) allow the approval decision to be reversed (i.e., recall) on the basis of post-approval information.[42] Costly misrepresentation of results can also be introduced by adding ex post lying costs à la Kartik, Ottaviani, and Squintani (2007)—another problem regulators are currently grappling with.[43] Future work could also add competition among researchers.[44] Finally, it would be interesting to extend our strategic analysis to settings with richer signal and state spaces, such as those pioneered by Moscarini and Smith (2001) and Fudenberg, Strack, and Strzalecki (forthcoming) respectively.

Times ("A Cancer Conundrum: Too Many Drug Trials, Too Few Patients," August 12, 2017) points out that there are currently 1000 immunotherapy trials underway, a rhythm that makes recruiting patients for clinical trials difficult.

[41] See also Taylor and Yildirim's (2011) static analysis.

[42] For example, see Zuckerman, Brown, and Nissen (2011) on the prevalence of recalls of drugs and medical devices. 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.

[43] Our 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). See also Shavell (1994), Dahm, Gonzalez, and Porteiro (2009), Henry (2009), Felgenhauer and Schulte (2014), and Herresthal (2017) on the incentives to disclose privately observed research results.

[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.

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

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

In this log-odds space, the lower and upper standards of research are $\left(\mathsf{s}{=}\ln\frac{s}{1-s}, \mathsf{S}{=}\ln\frac{S}{1-S}\right)$, the logit transformation 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 x_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 $x_t = x$ under the two states is

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

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

$$\mu' = \frac{2\mu^2}{\rho^2} \tag{3}$$

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 $\sigma \in (\mathsf{s}, \mathsf{S})$ we have that the expected payoff of player j (with cost of research c_j and benefits v_j^G and v_j^B from approval) is

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

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 conditional probabilities

$$\begin{aligned} \Psi(\sigma, \omega) &= E[e^{-rT} | \sigma(T) = \mathsf{S}, \omega] \ \Pr[\sigma(T) = \mathsf{S} | \omega] \\ \psi(\sigma, \omega) &= E[e^{-rT} | \sigma(T) = \mathsf{s}, \omega] \ \Pr[\sigma(T) = \mathsf{s} | \omega], \end{aligned} \tag{4}$$

we recover the expected discounted probabilities introduced in the main text

$$\begin{aligned}\Psi(\sigma) &= \Pr[\omega=G]\Psi(\sigma,G)+\Pr[\omega=B]\Psi(\sigma,B)\\ \psi(\sigma) &= \Pr[\omega=G]\psi(\sigma,G)+\Pr[\omega=B]\psi(\sigma,B).\end{aligned}$$

The expected payoff at $\sigma \in (\mathsf{s},\mathsf{S})$ is

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

The first line collects the research cost in perpetuity and the expected payoff if the upper standard S is reached first, taking into account the savings in future research costs. The second line is the expected payoff when the lower standard s is reached first. From Stokey (2009) we obtain the following closed-form expressions for the expected discounted conditional probabilities

$$\begin{array}{ll}\Psi(\sigma,G)=\frac{e^{-R_1(\sigma-\mathsf{s})}-e^{-R_2(\sigma-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}} & \Psi(\sigma,B)=\frac{e^{R_2(\sigma-\mathsf{s})}-e^{R_1(\sigma-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}\\ \psi(\sigma,B)=\frac{e^{-R_1(\mathsf{S}-\sigma)}-e^{-R_2(\mathsf{S}-\sigma)}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}} & \psi(\sigma,G)=\frac{e^{R_2(\mathsf{S}-\sigma)}-e^{R_1(\mathsf{S}-\sigma)}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}},\end{array} \tag{5}$$

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 expected discounted conditional probabilities Ψ and ψ satisfy*

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

where

$$\begin{array}{ll}a=\frac{R_1-R_2}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}<0 & b=\frac{R_2e^{R_2(\mathsf{S}-\mathsf{s})}-R_1e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}>0\\ f=\frac{R_1e^{-R_1(\mathsf{S}-\mathsf{s})}-R_2e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}<0 & g=\frac{R_2-R_1}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}>0.\end{array}$$

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:*

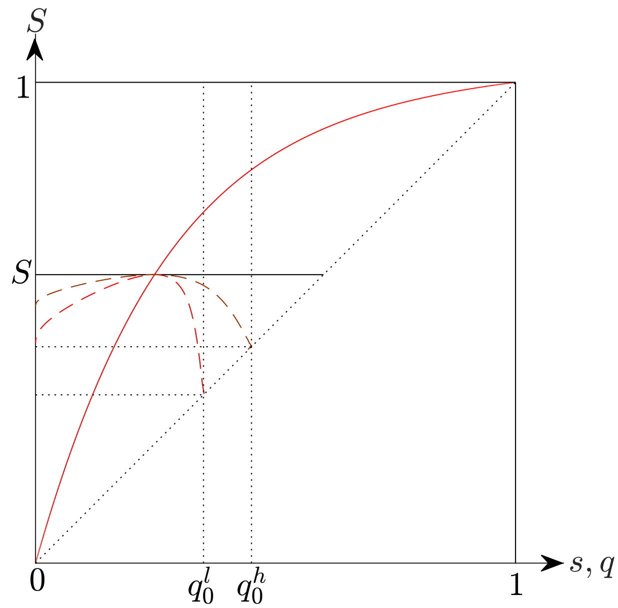

Figure 11: Best reply as initial belief changes.

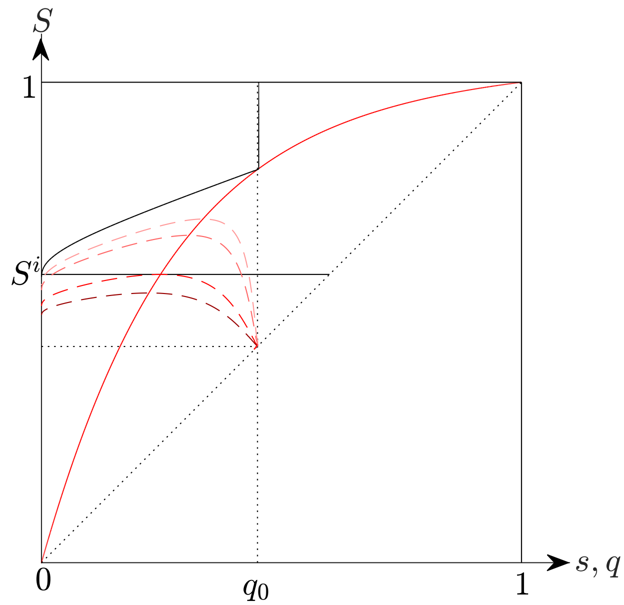

Figure 12: Construction of informer best reply.

***(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.*

Proof of Proposition 1

(a.i) This result follows as a special case of Lemma B2 with player j as the informer (with $v_j^G = v_j^B = v_i > 0$ and $c_j = c$). Figure 11 illustrates that the best reply does not depend on q. The figure plots isopayoff curves for $q = q_0^l$ and $q = q_0^h > q_0^l$. The lower best reply to S (corresponding to the tangency point with the horizontal line at S) is the same regardless of the value of q, even though the shape of the isopayoff curve and the corresponding payoff vary with q.

(a.ii) We prove that $b_i(S)$ is increasing in S, as illustrated in Figure 12. Using the log-odds parametrization and applying the implicit function theorem we have

$$\frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}} = -\left.\frac{\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{s} \partial \mathsf{S}}}{\frac{\partial^2 u_i(\sigma)}{\partial \mathsf{s}^2}}\right|_{\mathsf{s}=b_i(\mathsf{S})}. \tag{6}$$

In the case of the informer, $v_j^G = v_j^B = v_i > 0$ and $c_j = c$, so that equation (16) in Appendix B can be written as

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

so that, taking derivatives,

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

Equation (18) characterizes the lower best reply for $\mathsf{s} = b_i(\mathsf{S})$, so that b is defined by the implicit function $\Upsilon(\mathsf{S}, b) = 0$ where

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

Taking the total differential we can compute

$$\frac{\partial b}{\partial \mathsf{S}} = -\frac{\left[\frac{\partial a}{\partial \mathsf{S}}(1+e^{-\mathsf{S}}) - ae^{-\mathsf{S}}\right]\left(v_i + \frac{c}{r}\right) + \frac{c}{r}\frac{\partial b}{\partial \mathsf{S}}(1+e^{-\mathsf{s}})}{\frac{c}{r}(1+e^{-\mathsf{s}})} = -\frac{\left[\frac{\partial a}{\partial \mathsf{S}}(1+e^{-\mathsf{S}}) - ae^{-\mathsf{S}}\right]\left(v_i + \frac{c}{r}\right)}{2\frac{c}{r}(1+e^{-\mathsf{s}})}.$$

Substituting in the numerator of expression (6), we have

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

which 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 (19) 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 $q \geq \hat{q}_e$, in an MPE the informer must choose $\mathscr{A}_i$.

For $q < \hat{q}_e$ we have that, if in equilibrium the informer chooses $\mathscr{I}_i$, 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 $V_i^i(q_0)$, study its limit when c and r converge to 0, and then compare it to the value obtained by KG.[45] 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 B2 (equation 18), we see that, as c converges to 0, s goes to $-\infty$.[46] In the limit with costless 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 expected payoff of the informer in this case is given by expression (12) in Appendix B. Using the fact that under informer authority $\mathsf{S}^i = \hat{\sigma}_e$, we have that the limit of $V_i^i(\sigma_0)$ as c converges 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 (5) 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 $V_i^i(\sigma_0)$, 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.i) This result follows as a special case of Lemma B1 with player j as the evaluator (i.e., $v_j^G = v_e^G$, $v_j^B = v_e^B$ and $c_j = 0$).

(a.ii) For the evaluator, 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

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

Expression (13) for the evaluator gives

$$\frac{\partial u_e(\sigma)}{\partial \mathsf{S}} = \frac{e^{\sigma}}{1+e^{\sigma}}\Psi(\sigma, G)\left[f\cdot\left(v_e^G + e^{-\mathsf{S}}v_e^B\right) - e^{-\mathsf{S}}v_e^B\right], \tag{8}$$

which implies that

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

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

[46] Equation (18) can be reexpressed as $e^{-\mathsf{s}}[b-1] = -a\left(1+e^{-\mathsf{S}}\right)\left(\frac{rv_i}{c}+1\right) - b$. Taking the limit as c converges to 0 yields the result.

where the second expression is derived from equation (14) characterizing $B_e(s)$. The numerator of (7) 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 of (7) is negative, as shown in (15) 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 $q = \overline{q}$. As a reminder, we consider MPE 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: *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 informer does not choose strategy $\mathscr{I}_i$, neither in substage 1 nor in substage 3. At both $q = \underline{q}$ and $q = \overline{q}$, the informer is thus called to decide $\mathscr{R}_i$ in substage 3 and between $\mathscr{A}_i$ and $\mathscr{R}_i$ in substage 1. By definition of the endpoints, if $q = \underline{q}$ or $q = \overline{q}$, the evaluator knows in substage 2 that the informer will choose $\mathscr{R}_i$ in substage 3. Thus in substage 2, if $\underline{q} > \hat{q}_e$ the evaluator chooses $\mathscr{A}_e$, while if $\underline{q} \leq \hat{q}_e$, the evaluator chooses $\mathscr{R}_e$. Hence, if $\underline{q} > \hat{q}_e$, the informer's best reply is to choose $\mathscr{A}_i$ in substage 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 $\mathscr{A}_e$ in substage 2. Indeed, $\mathscr{R}_e$ 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 2: *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 1 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 $\underline{q} = b_i(\overline{q})$ and $\overline{q} = B_e(\underline{q})$. For any belief $q < \hat{q}_e$ in the interval $(\underline{q}, \overline{q})$, if the informer chooses $\mathscr{A}_i$, the evaluator chooses $\mathscr{R}_e$ and the informer responds by choosing $\mathscr{I}_i$, because research is the equilibrium outcome for $q \in (\underline{q}, \overline{q})$. The informer is thus indifferent between choosing $\mathscr{A}_i$ or $\mathscr{I}_i$ in substage 1. 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 $\mathscr{R}_e$ (that is followed by $\mathscr{I}_i$ by the informer). By the same logic and by definition of B_e, it has to be the case that $\overline{q} = B_e(\underline{q})$.

Step 3: *MPE exists and is unique.*

We established that the equilibrium outcome of the no-commitment game, if it exists, is at the intersection of $b_i(S)$ and $B_e(s)$. Also, $b_i(S)$ is increasing in S and $B_e(s)$ 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.[47] 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 $\max_S u_e|_{s=b_i(S)}$. The solution to this problem, $S^e(q_0)$, depends on the starting belief q_0.

Step 1: *For any q_0, $S^i \leq S^e$ and $s^i \leq s^e$.*

$S^e(q_0)$ 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. Furthermore, given that $b_i(S)$ 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.

Step 2: *For any q_0, $S^e \leq S^n$ and $s^e \leq s^n$.*

Notice that

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

From the definition of S^n we have that $\frac{\partial u_e}{\partial \mathsf{S}}(b_i(\mathsf{S}),\mathsf{S}) \leq 0$ for $S \geq S^n$. Replacing in expression (9), and using the fact that u_e is decreasing in s and $b_i(\mathsf{S})$ is increasing in S, we obtain $\left.\frac{\partial u_e}{\partial \mathsf{S}}\right|_{\mathsf{s}=b_i(\mathsf{S})} < 0$ at

[47] 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.

$S \geq S^n$. Thus, we have $S^e \leq S^n$. Moreover, since $b_i(S)$ is increasing in S, it follows that $b_i(S^e(q_0)) \leq b_i(S^n) = s^n$ for any q_0. Given these general properties, we distinguish three cases:

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

(b) $q_0 \in \left(s^i, S^n\right)$. 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

$$\frac{\partial S^e}{\partial \sigma_0} = - \left. \frac{\frac{\partial^2 u_e(\sigma_0)}{\partial \mathsf{S} \partial \sigma_0}}{\frac{\partial^2 u_e(\sigma_0)}{\partial \mathsf{S}^2}} \right|_{\mathsf{s}=b_i(\mathsf{S})} . \tag{10}$$

Since b_i is independent of σ_0 by Proposition 1.a.i, using expression (9), we have

$$\left. \frac{\partial^2 u_e}{\partial \mathsf{S} \partial \sigma_0} \right|_{\mathsf{s}=b_i(\mathsf{S})} = \frac{\partial^2 u_e}{\partial \mathsf{s} \partial \sigma_0} (b_i(\mathsf{S}), \mathsf{S}) \, \frac{\partial b_i(\mathsf{S})}{\partial \mathsf{S}} + \frac{\partial^2 u_e}{\partial \mathsf{S} \partial \sigma_0} (b_i(\mathsf{S}), \mathsf{S}) . \tag{11}$$

The first term is positive because its first factor is positive by Lemma C1 in Appendix C and its second factor is also positive by Proposition 1.a.ii. The second term is positive by Lemma C2. Finally, the denominator of (10) is negative by Lemma C3, proving that (10) is positive.

(c) $q_0 \in (S^n, 1)$. 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) 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 $b_w(S)$ always 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 $s_w^* = b_w(S_w^*) \leq s^i = b_w(\hat{q}_e)$. Overall we have $s_w^* \leq s^i \leq s^e \leq s^n$.

(b) The approval standard under informer authority $S^i = \hat{q}_e$ is independent of c. The socially optimal S_w^* is decreasing in c. Furthermore in the limit when c is large, $\lim_{c \to +\infty} S_w^* = \hat{q}_w < \hat{q}_e$.

If r is sufficiently low then $S_w^* > \hat{q}_e$ at $c = 0$ and therefore we can find $\bar{c} > 0$ such that if $c = \bar{c}$, $S_w^* = S^i = \hat{q}_e$. If instead r is sufficiently high then $S_w^* < \hat{q}_e$ at $c = 0$ and therefore $\bar{c} = 0$.

(i) If $c > \bar{c}$, $S_w^* < S^i$ follows from $S^i \leq S^e \leq S^n$ as argued in the main text.

(ii) If $c < \bar{c}$, we have $S_w^* > S^i$. Furthermore we established in Section 1.3 that the evaluator commitment is an increasing function of q_0 ranging from S^i to S^n. Since $S^n > S_w^*$ and S_w^* is independent of q_0, we can identify $\overline{q}_0(c)$ such that, if the initial belief is $q_0 = \overline{q}_0(c)$, $S^e = S_w^*$. Given S^e is increasing in q_0 this implies that $S^i \leq S^e \leq S_w^* \leq S^n$ when $q_0 < \overline{q}_0(c)$ and that $S^i \leq S_w^* \leq S^e \leq S^n$ when $q_0 > \overline{q}_0(c)$. Notice that if $\bar{c} = 0$ **(i)** holds for all c.

Proof of Proposition 5

(a) 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.

(b) As indicated in Proposition 4, for $c > \bar{c}$, informer authority dominates evaluator commitment both in terms of false positive and false negative, and thus informer authority achieves a strictly higher social welfare level. For $c = \bar{c}$, informer authority achieves the first best and strictly dominates evaluator commitment. By continuity, there exists $\tilde{c} < \bar{c}$ such that for $c > \tilde{c}$, informer authority welfare dominates evaluator commitment.

Proof of Proposition 6

A key property used in this proof and visible in Figure 13 is that the informer-authority 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 raised, $\hat{q}_e(P)$ increases and the informer-authority outcome (s^i, S^i) moves up along the planner's lower best reply curve b_w. In particular, since the socially optimal outcome is independent of P, there is a boundary price $\tilde{P}$ at which the informer-authority outcome coincides with the socially optimal solution. For any $P > \tilde{P}$, informer authority dominates evaluator commitment.

Proof of Proposition 7

(i) $q_0 \in (s^I, \hat{q}_w)$. In the game with the planner, costly research is undertaken given that $s^I < q_0$, but this research is socially worthless because $S^I = \hat{q}_w$; thus the planner obtains a strictly negative payoff by playing directly. If the role of player is delegated to the evaluator, when $q_0 \in (s^I, s^i]$ 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

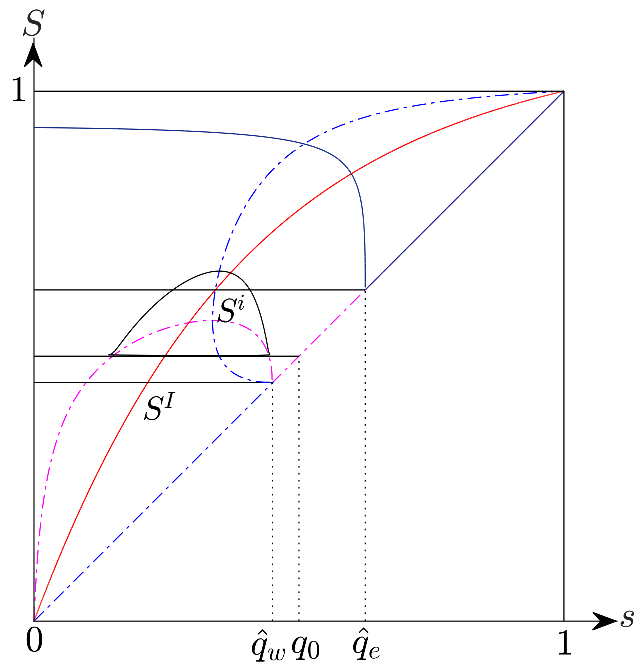

Figure 13: Value of delegation.

under informer authority playing against the evaluator 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 replying 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. 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 isopayoff 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 13. The planner's payoff at

(s^i, S^i) 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 (s^i, S^i) 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 isopayoff lens corresponding to the approval payoff at q_0 collapses into a point, moreover there exists $\underline{S}$ such that $b_w(q_0) = b_w(\underline{S})$ and we can find $\check{P}(q_0)$ such that $S^i = \underline{S}$ and $\hat{P}(q_0)$ such that $S^i = q_0$. Since $\hat{P}(q_0) = q_0$ then for any $P \in (\check{P}(q_0), \hat{P}(q_0))$ the outcome is always immediate approval regardless of whether the planner or the evaluator is playing. If $P > \hat{P}(q_0)$ then the planner is strictly better off taking up the role of player, indeed any outcome with research yields the planner a lower payoff.

Proof of Proposition 8

(a) The planner can always block research by committing to an upper cutoff $S^W(q_0)$ above $b_i^{-1}(q_0)$, inducing the informer to set the lower standard $s^W(q_0)$ 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 $q_0 \leq s_w^*$ 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, $\tilde{q} \leq q_0 < S^N$, the planner optimally chooses an interior commitment, $S^W(q_0)$. At $\tilde{q}$, the planner is indifferent between rejecting the project, thus blocking research, and committing to an interior benchmark $S^W(\tilde{q}) = b_i^{-1}(\tilde{q})$ 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 $S^W(q_0)$ compares to the Nash standard, S^N, 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 to the left the planner's lower best reply, and conversely if $S > \check{S}$.

(i) When $P < P^*$, as in Figure 14, $\tilde{q}$ 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 $\tilde{q} \leq q_0 < S^N$, the planner commits to $S^W(q_0) < S^N$. Moreover, as $q_0 \to S^N$, the commitment solution moves along an upward-sloping path that converges to the unique Nash outcome ($S^W(q_0) \to S^N$). For a belief greater than

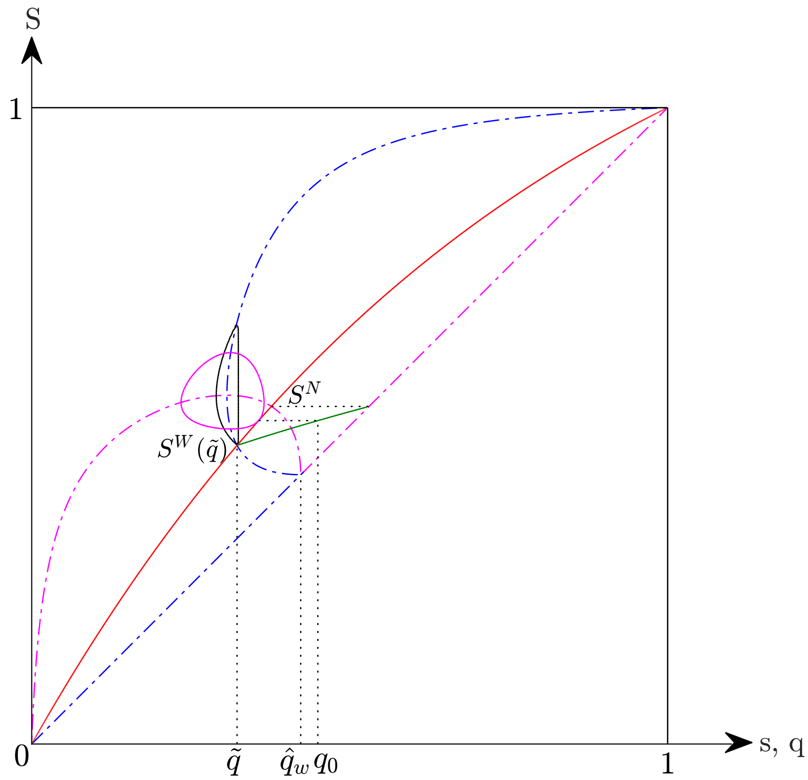

Figure 14: Increasing commitment path.

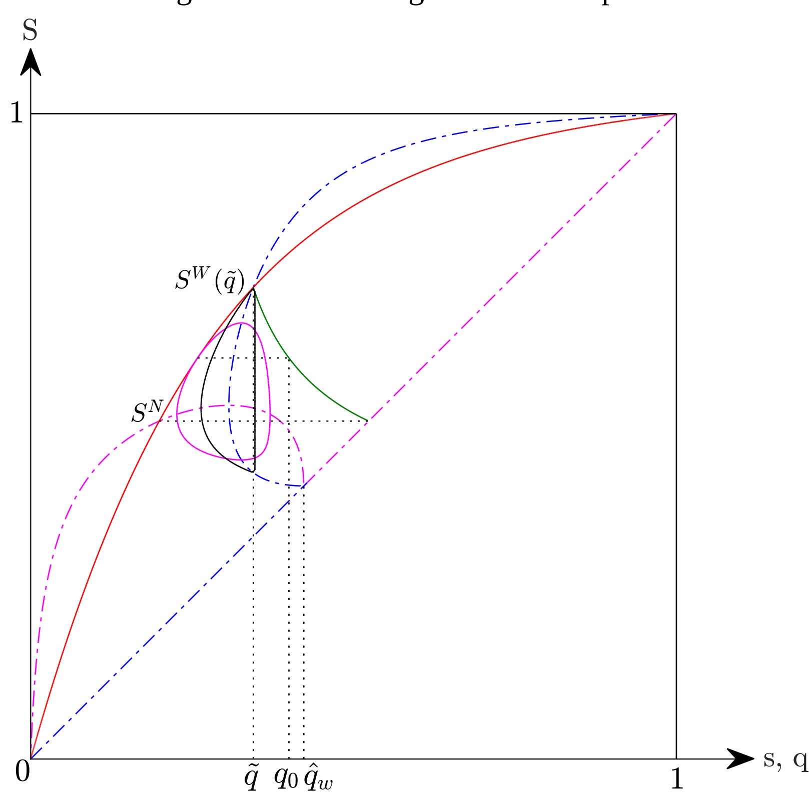

Figure 15: Decreasing commitment path.

$\tilde{q}$, in fact, the isopayoff curve of the planner (purple in Figure 14) is tangent to the informer's lower best reply for $s^i > \tilde{q}$.

For $P = P^*$, 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 $S^W(q_0) = S^N$ for any initial belief $q_0 \in (\tilde{q}, S^N)$.

(**ii**) When $P^* < P \leq \underline{P}$, as in Figure 15, $\tilde{q}$ 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 $\tilde{q} \leq q_0 < S^N$, the optimal commitment is above the Nash level, $S^W(q_0) > S^N$. In contrast with the decision-maker commitment, as $q_0 \to S^N$, the planner's commitment solution moves along a downward (rather than upward) sloping path that converges to the Nash equilibrium of the game, $S^W(q_0) \to S^N$. Graphically, the tangency point is now on the left of $\tilde{q}$, as the purple indifference curve in Figure 15 highlights.

(**iii**) For P slightly above $\underline{P}$, the informer's lower best reply (red curve in Figure 9) crosses the planner's zero-level isopayoff 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 isopayoff (continuous black) is tangent to the informer's lower best reply at a belief, $b_i(S^W(\tilde{q}))$, 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 $S = S^W(\tilde{q})$ commitment. Since $b_i(S^W(\tilde{q})) < \tilde{q}$, such an interior commitment induces the informer to set the lower standard below $\tilde{q}$. The planner's optimal interior commitment $S^W(\tilde{q})$ 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.

PARAMETERS	μ	v_i	c	r	v_e^G	v_e^B	
Figures 1 and 12	12	1.7	15	5	1	-1.5	$q_0 = 0.45$
Figure 2: continuous curve	12	1.7	0	0	1	-1.5	$q_0 = 0.45$
Figure 2: dashed curve	12	1.7	0	5	1	-1.5	$q_0 = 0.45$
Figure 2: dotted curve	12	1.7	15	5	1	-1.5	$q_0 = 0.45$
Figures 3 and 4	12	1.7	15	5	0.5	-0.5	$q_0 = 0.45$
Figure 5	12	1.7	15	5	0.5	-0.5	$q_0 = 0.45$, $q_0' = 0.65$
Figure 6	5	1.5	15	0.8	4	-6	$q_0 = 0.55$
Figure 7	5	1.5	3	0.8	4	-6	$\overline{q}_0 = 0.36$
Figure 8	5	1.5	10	0.8	4	-6	$q_0 = 0.55$
Figures 9 and 10	5	0.5	15	0.8	7	-6	$P = 5$
Figure 11	12	1.7	15	5	1	-1.5	$q_0^l = 0.35$, $q_0^h = 0.45$
Figure 13	5	0.6	15	0.8	6	-6	$P = 1.5$
Figure 14	5	0.5	15	0.8	7	-6	$P = 0.1$
Figure 15	5	0.5	15	0.8	7	-6	$P = 3.5$

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.

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*, 84(3): 969–1014.

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

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.

Carpenter, Daniel. 2004. “Protection without Capture: Product Approval by a Politically Responsive, Learning Regulator.” *American Political Science Review*, 98(4): 613–631.

Carpenter, Daniel. 2011. *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*. Princeton and Oxford: Princeton University Press.

Carpenter, Daniel, and Michael M. Ting. 2007. “Regulatory Errors with Endogenous Agendas.” *American Journal of Political Science*, 51(4): 835–852.

Chan, Jimmy, Alessandro Lizzeri, Wing Suen, and Leeat Yariv. 2018. “Deliberating Collective Decisions.” *Review of Economic Studies*, 85(2): 929–963.

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. Forthcoming. "Speed, Accuracy, and the Optimal Timing of Choices." *American Economic Review*.

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.

Henry, Emeric, Marco Loseto, and Marco Ottaviani. 2018. "Regulation with Experimentation: Ex Ante Approval, Ex Post Withdrawal, and Liability." Sciences Po and Bocconi University.

Herresthal, Claudia. 2017. "Hidden Testing and Selective Disclosure of Evidence." Queen's College, University of Cambridge.

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*.

Marks, Harry M. 1997. *The Progress of Experiment: Science and Therapeutic Reform in the United States, 1900–1990*. New York: Cambridge University Press.

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

McClellan, Andrew. 2017. "Experimentation and Approval Mechanisms." Unpublished working paper.

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.

Morris, Stephen, and Philipp Strack. 2017. "The Wald Problem and the Equivalence of Sequential Sampling and Static Information Costs." SSRN Working Paper.

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." SSRN Working Paper.

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. 1945. "Sequential Tests of Statistical Hypotheses." *Annals of Mathematical Statistics*, 16(2): 117–186.

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.

B Online 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) First-Order Condition for the Upper Best Reply. By parts (1) and (2) of Lemma B0 player j's expected payoff $u_j(\sigma)$ can be written as

$$u_j(\sigma)=-\frac{c_j}{r}+\frac{e^{\sigma}\Psi(\sigma,G)}{1+e^{\sigma}}\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]+\frac{e^{\sigma}}{1+e^{\sigma}}\psi(\sigma,G)(1+e^{-\mathsf{s}})\frac{c_j}{r}. \tag{12}$$

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

$$\frac{\partial u_j(\sigma)}{\partial\mathsf{S}}=\frac{e^{\sigma}\Psi(\sigma,G)}{1+e^{\sigma}}\left\{f\cdot\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]-e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)+g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}\right\}, \tag{13}$$

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

$$f\cdot\left[v_j^G+e^{-\mathsf{S}}v_j^B+\left(1+e^{-\mathsf{S}}\right)\frac{c_j}{r}\right]=e^{-\mathsf{S}}\left(v_j^B+\frac{c_j}{r}\right)-g\cdot(1+e^{-\mathsf{s}})\frac{c_j}{r}. \tag{14}$$

Equation (14) 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 $c_e=0$, (14) simplifies into $v_e^G+e^{-\mathsf{S}}v_e^B=\frac{e^{-\mathsf{S}}v_e^B}{f}$.

Second-Order Condition for the Upper Best Reply. Differentiating (13) with respect to S we have

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

Equation (14) then implies

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

Some algebra yields

$$\begin{aligned}1-f&=\frac{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}-R_1e^{-R_1(\mathsf{S}-\mathsf{s})}+R_2e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}\\ &=\frac{R_2e^{-R_1(\mathsf{S}-\mathsf{s})}-R_1e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}=\frac{R_2e^{R_2(\mathsf{S}-\mathsf{s})}-R_1e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}=-\frac{\partial g}{\partial\mathsf{S}}\frac{1}{g}.\end{aligned}$$

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

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

which, by equation (14), can be rewritten as

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

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 $\mathsf{S}=B_j(\mathsf{s})$, we conclude that

$$\left.\frac{\partial^2 u(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}=B_j(\mathsf{s})}<0 \tag{15}$$

if and only if $\frac{\partial f}{\partial \mathsf{S}}<-f(1-f)$, i.e.,

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

which always holds being equivalent to $2e^{-(\mathsf{S}-\mathsf{s})}<e^{-2R_1(\mathsf{S}-\mathsf{s})}+e^{-2R_2(\mathsf{S}-\mathsf{s})}\Leftrightarrow 0<(e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})})^2$.

(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

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

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

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

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 $\tilde{\mathsf{S}}$ such that $\left.\frac{\partial u_j(\sigma)}{\partial \mathsf{S}}\right|_{\mathsf{S}=\tilde{\mathsf{S}}} \geq 0$. 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 an interior solution $\mathsf{S}^* \leq \tilde{\mathsf{S}}$ to $\frac{\partial u_j(\sigma)}{\partial \mathsf{S}} = 0$ such that $\left.\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{S}^2}\right|_{\mathsf{S}^*=B_j(\mathsf{s})} \geq 0$, 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) First-Order Condition for the Lower Best Reply. By parts (3) and (4) of Lemma B0**,** taking a derivative of (12) with respect to s yields

$$\frac{\partial u_j(\sigma)}{\partial \mathsf{s}} = \frac{e^\sigma \psi(\sigma, G)}{1+e^\sigma} \left\{ a \left[v_j^G + e^{-\mathsf{S}} v_j^B + \left(1+e^{-\mathsf{S}}\right) \frac{c_j}{r} \right] + \frac{c_j}{r} \left[b(1+e^{-\mathsf{s}}) - e^{-\mathsf{s}} \right] \right\}. \tag{16}$$

Hence, player j's first order condition is

$$v_j^G + e^{-\mathsf{S}} v_j^B + \left(1+e^{-\mathsf{S}}\right) \frac{c_j}{r} = -\frac{1}{a} \frac{c_j}{r} \left[b(1+e^{-\mathsf{s}}) - e^{-\mathsf{s}} \right] \tag{17}$$

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 (17) simplifies into

$$a\left(1+e^{-\mathsf{S}}\right)\left(v_i+\frac{c}{r}\right)+\frac{c}{r}\left[b(1+e^{-\mathsf{s}})-e^{-\mathsf{s}}\right]=0. \tag{18}$$

Second Order Condition for the Lower Best Reply. Taking a derivative with respect to s of (16) 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 (17), 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

$$\begin{aligned}
1-b &= \frac{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}-R_2e^{R_2(\mathsf{S}-\mathsf{s})}+R_1e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}\\
&= \frac{R_1e^{R_2(\mathsf{S}-\mathsf{s})}-R_2e^{R_1(\mathsf{S}-\mathsf{s})}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}=\frac{R_1e^{-R_1(\mathsf{S}-\mathsf{s})}-R_2e^{-R_2(\mathsf{S}-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}=-\frac{\partial a}{\partial \mathsf{s}}\frac{1}{a},
\end{aligned}$$

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 $\frac{\partial b}{\partial \mathsf{s}}<-b(1-b)$, i.e.,

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

which always holds being equivalent to $2e^{(\mathsf{S}-\mathsf{s})}<e^{2R_1(\mathsf{S}-\mathsf{s})}+e^{2R_2(\mathsf{S}-\mathsf{s})}$. Thus,

$$\left.\frac{\partial^2 u_j(\sigma)}{\partial \mathsf{s}^2}\right|_{\mathsf{s}=b_j(\mathsf{S})}<0. \tag{19}$$

(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

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

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,

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

Focusing on the second factor, we obtain

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

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 $\mathsf{S}>\hat{\sigma}_j$. 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 $\mathsf{S}\le\hat{\sigma}_j$. In this case we show that $\frac{\partial u(\sigma)}{\partial\mathsf{s}}>0$. To see this, assume by contradiction that there exists $\tilde{\mathsf{s}}$ such that $\left.\frac{\partial u_j(\sigma)}{\partial\mathsf{s}}\right|_{\mathsf{s}=\tilde{s}}\le 0$. 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 an interior solution $\mathsf{s}^*\ge\tilde{\mathsf{s}}$ to $\frac{\partial u_j(\sigma)}{\partial\mathsf{s}}=0$ such that $\left.\frac{\partial^2 u_j(\sigma)}{\partial\mathsf{s}^2}\right|_{\mathsf{s}^*=b(\mathsf{S})}\ge 0$, 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 Online Supplementary Appendix C: Technical Results

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

$$\frac{\partial^2 u_e}{\partial \mathsf{s} \partial \sigma} > 0. \tag{20}$$

Proof of Lemma C1
Using equation (16) from Appendix B for $c_j = 0$ we have

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

so that, since a does not depend on σ,

$$\frac{\partial^2 u_e}{\partial \mathsf{s} \partial \sigma} = \frac{\partial \left(\frac{e^{\sigma}}{1+e^{\sigma}} \psi(\sigma, G) \right)}{\partial \sigma} a \left[v_e^G + e^{-\mathsf{S}} v_e^B \right]. \tag{21}$$

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

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

Overall, replacing in equation (21), and using $a < 0$, we obtain (20).

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

$$\left. \frac{\partial^2 u_e}{\partial \mathsf{S} \partial \sigma} \right|_{\mathsf{s} = b_i(\mathsf{S})} > 0. \tag{22}$$

Proof of Lemma C2
Using (8) from Appendix B we have

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

given that f is independent of σ. Thus,

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

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 (22).

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

$$\left.\frac{\partial^2 u_e}{\partial S^2}\right|_{s=b_i(S)} < 0 \text{ for } S \leq S^n. \tag{23}$$

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

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

Using the expression for the evaluator's expected payoff (12) for $c_j = 0$ and $j = e$, we now show that the four terms in (24) are negative so that we have (23):

- 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 evaluator's expected payoff 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, (23) follows.