Regulation with Experimentation: Ex Ante Approval, Ex Post Withdrawal, and Liability*

Emeric Henry[†] Marco Loseto[‡] Marco Ottaviani[§]

June 23, 2021

Abstract

We analyze the optimal mix of ex ante experimentation and ex post learning for the dynamic adoption of activities with uncertain payoffs in a two-phase model of information diffusion. In a first pre-introduction phase, costly experimentation is undertaken to decide whether to adopt an activity or abandon experimentation. In a second stage following adoption, learning can continue possibly at a different pace while the activity remains in place; the withdrawal option is exercised following the accumulation of sufficiently bad news. We compare from a law and economics perspective the performance of three regulatory frameworks commonly adopted to govern private experimentation and adoption incentives: liability, withdrawal, and authorization regulation. Liability should be preempted to avoid chilling of activities that generate large positive externalities, consistent with the preemption doctrine. Liability should be used to discourage excessive experimentation for activities that generate small positive externalities. Authorization regulation should be lenient whenever it is used, consistent with the organization of regulation in a number of areas ranging from product safety to antitrust.

Keywords: Authorization regulation, liability, withdrawal, experimentation, preemption doctrine.

JEL Classification: D18 (Consumer Protection), D83 (Learning; Information and Knowledge), K13 (Product Liability), K2 (Regulated Industries), M38 (Government Policy and Regulation).

---

*We thank Florian Baumann, Jean-Pierre Benoît, Alessandro Bonatti, Jaroslav Borovička, Daniel Carpenter, Andrew Daughety, Francesco Decarolis, Maryam Farboodi, Lars P. Hansen, Xinyu Hua, Stefano Pegoraro, Federico Pessina, Nicolas Schutz, Kathryn Spier, and seminar participants at Joint Program Conference 2021 (Chicago), Economic Dynamics workshop 2020 (Chicago), ASSA-ES 2019 (Atlanta), Bonn, CEPR Applied IO 2019 (Madrid), EALE 2018 (Milan), Ecole Polytechnique, EEA 2019 (Manchester), ESSET 2017 (Gerzensee), LBS, and UCL for helpful comments.

[†]Sciences Po and CEPR, emeric.henry@sciencespo.fr

[‡]University of Chicago, mloseto@uchicago.edu

[§]Bocconi University, BIDSA, CEPR, and IGIER, marco.ottaviani@unibocconi.it

1 Introduction

The COVID-19 crisis has turned the spotlight on the question of drug regulation. The need for a vaccine has fueled the debate between those who favor ex ante regulation relying on rigorous randomized control trials versus the camp supporting faster approval and the use of ex post regulation based on real-world evidence (Angus 2020). Beyond drugs, frameworks for regulating the introduction of risky products vary across industries. While the average delay from pre-clinical testing to FDA approval is twelve years for drugs and three to seven years for medical devices (Van Norman 2016), the U.S. Federal Aviation Administration (FAA) authorized the Boeing 737 MAX following a series of tests that lasted less than two years. Why is it easier for an aircraft manufacturer to have an airplane certified to fly than for a drug company to have a drug approved? For airplane safety regulation the threat of significant liability (post market introduction) seems to play a more systematic role than the stringency of ex ante approval (pre market introduction).

Meanwhile in digital markets, as competition concerns have been mounting, the opportunity of shifting to ex ante regulation is gaining momentum. Scott-Morton et al.'s (2019) Stigler Report and the Furman Review (2019) propose the creation of a digital agency that would set ex ante rules to regulate potentially anti-competitive and harmful practices by firms that dominate these markets. The new rules would serve as a substitute for ex post antitrust enforcement, where liability currently plays a key role especially in the US. As stated in the Furman Review, the idea is to move "from a purely ex post approach towards ex ante monitoring and enforcement of a clearer and more detailed set of rules."[1]

This paper analyzes the law and economics of regulation of private activities with uncertain social returns, such as new drugs or business practices in digital markets. We focus on the following three regulatory frameworks—including both ex ante and ex post interventions—that evolved over the course of the twentieth century, as decision rights were gradually transferred from judges to modern-day regulatory agencies (see Glaeser and Shleifer, 2003):

- *Liability* was the first form of regulation to appear, mostly through litigation. The threat of liability affects ex post withdrawal decisions by firms, but also shapes ex ante incentives for product introduction.
- *Withdrawal regulation*, granting the power to order withdrawal of unsafe products from the market, was the next regulatory tool to emerge. Like liability, withdrawal is also an ex post instrument, but shapes ex ante decisions differently.

[1] See Cabral (2021) for further discussion.

- *Authorization regulation* eventually appeared with the establishment of regulatory agencies (FDA for drugs in 1906, FAA for airline safety in 1958, and EPA for environmental projects in 1970). Agencies were empowered to authorize market introduction as well as to order ex post withdrawal following the accumulation of sufficiently bad news.

To compare the performance of these regulatory frameworks from a positive and normative perspective, we formulate a novel and tractable two-stage model with ex ante experimentation and ex post learning. First, a phase of ex ante experimentation with market research and testing takes place prior to market introduction. This initial experimentation is carried out in a rather controlled environment, for example through carefully monitored randomized controlled trials.[2] In the ex ante experimentation phase, the decision is eventually taken to abandon experimentation or to enter the market. Second, in the ex post learning phase after market introduction, learning continues often in a less controlled way and at a different pace, possibly leading to a final withdrawal decision following the accumulation of sufficiently bad news.

Section 2 sets the stage by formulating the decision-theoretic first-best benchmark (planner solution) of our continuous-time model with discounting. Public information is accumulated in the form of a Wiener process whose drift depends on a binary state of the world, either good or bad, corresponding to the social desirability of the activity. In the first (ex ante) phase, information arrival results from costly experimentation. Decision payoffs are collected only in the second (ex post) stage after the activity is implemented; the payoff of the planner is positive in the good state $v_p^G > 0$ and negative in the bad state $v_p^B < 0$. In the second stage, information arrives at a possibly different speed from the first stage, but now with no direct cost.

As shown in Section 2.1, the planner solution strikes an optimal balance between ex ante testing and ex post surveillance:

- The ex post phase corresponds to a bandit problem, where the planner chooses between a safe arm (withdraw) and a risky arm (continue undertaking the activity). If the belief q about the state being good—the state variable representing the posterior belief summarizing all information at each point in time—becomes sufficiently low, $q \leq z^*$, the planner withdraws.
- The ex ante phase is a generalized Wald (1945) problem, characterized by an experimentation cutoff s^* and an adoption cutoff S^*.[3] The planner abandons (rejects) the project as soon as the belief q falls below s^*, and implements the activity (adopts) as soon as the belief reaches S^*. The social planner collects information for beliefs between these two cutoffs $q \in (s^*, S^*)$.

[2]More generally, theoretical research and "dry bench" experiments have similar features.

[3]This is not a standard Wald problem since the adoption payoff is state dependent.

In practice, firms control the experimentation process. The incentives of the firm are typically misaligned with the planner. In particular, the firm does not suffer the full social damage in the bad state and does not necessarily recoup the full social benefits in the good state. Absent regulation, the payoff collected by the firm in the second phase is v_f, independent of the state, with $0 < v_f \leq v_p^G$. The firm thus generates a positive externality $e^G = v_p^G - v_f$ on the rest of society in the good state, a key parameter in our analysis. If unregulated, the firm would immediately introduce the product and never withdraw.

The role of regulation is to align private and social incentives by shaping the experimentation (s), adoption (S), and withdrawal (z) decisions. The second part of the paper analyzes the effectiveness of three regulatory forms used to control these decisions: liability (with penalties imposed in case the state is bad), ex post withdrawal (ordered by a judge or regulator), and authorization regulation (whereby a regulator ex ante approves market introduction and then controls ex post withdrawal).

Section 3.2 examines liability, understood as a liability rate charged per unit of time of market presence in the ex post stage when the state is bad, so that the flow payoff of the firm in the bad state becomes $v_f - L$. This liability is strict, i.e., it does not depend on proof of negligence or intent. Without externality in the good state ($e^G = 0$), the first best is achieved by a liability rule requiring the firm to fully compensate for the whole social damage in the bad state, because incentives are then perfectly aligned in both states. However, when the externality is positive ($e^G > 0$), the planner should commit to being more lenient than under full liability.

Adjusting the liability L, the planner can induce the firm to withdraw at any standard z.[4] In particular the planner could choose the liability rate $\hat{L}$ that induces first-best withdrawal z^*. This requires the payoff of the firm in the bad state to be decreased to $\hat{L}$ in such a way that the ratio of private to social payoffs is the same in both states. Thus, when liability is set at $\hat{L}$, the firm obtains lower profits than the planner in both states. As a consequence, the firm has lower incentives to experiment and withdraws too early compared to the planner. The firm also attaches less value to information and thus introduces too early. Moving away from $\hat{L}$ clearly increases welfare—the first-order gain in ex ante incentives trumps the ex post loss (of second order given that $\hat{L}$ induces the optimal withdrawal).[5]

[4]Garber (2013) presents evidence that product liability resulted in the withdrawal of the drug Benedictin in 1983 as well as of a number of vaccines during the 1980s and 1990s.

[5]Whether the optimal liability should be set above or below $\hat{L}$ depends on the tradeoff between the conflicting effects on experimentation and approval. If the initial belief q is low, encouraging experimentation is most pressing, so it is optimal for the planner to be more lenient. If q is high, the planner should instead be tougher to discourage early approval.

Section 3.3 turns to withdrawal regulation, understood as a commitment to withdraw the first time the belief reaches standard z. To appreciate the difference between liability and withdrawal regulation, set the liability at a level that induces the firm to withdraw at the same z. In this case, from the ex post perspective liability and withdrawal are equivalent by design. As we show, however, these tools induce different ex post value functions and thus result in different ex ante incentives, for two reasons. First, liability is more taxing and thus reduces the value function of the firm compared to withdrawal regulation. This force decreases the option value of experimentation, for any fixed adoption standard S, thus pushing the firm to experiment less under liability. Second, while under liability the value is tangent at the point of voluntary withdrawal (by the smooth pasting optimality condition for the firm), the ex post value function under withdrawal regulation features a kink at z. Because of this kink, under withdrawal regulation, the firm has an incentive to delay approval to postpone a potential withdrawal and increase the time on the market. If the cost of information acquisition is sufficiently low, we show that this postponement effect dominates, thus leading to later adoption under withdrawal regulation than under liability (resulting in the same withdrawal level).

Section 3.4 analyzes authorization regulation whereby the agency controls not only ex post withdrawal but also ex ante approval, while the firm still controls ex ante experimentation. Suppose that both approval and withdrawal are set at the first-best levels. When $e^G = 0$, the firm experiments too much, because it obtains the social payoff in the good state but does not incur the loss in the bad state. When instead e^G is sufficiently high, the firm has insufficient incentives to experiment. As we show, there exists an intermediate value of the externality $\hat{e}$ at which the firm chooses precisely the socially optimal level of experimentation. When the externality is lower, $e^G < \hat{e}$, it is optimal for the planner to discourage experimentation by setting tougher approval *and* withdrawal standards compared to the first best. If instead $e^G > \hat{e}$, the planner encourages experimentation by being more lenient, again in terms of both approval and withdrawal.

Section 3.5 derives policy implications by characterizing the socially optimal mix of the three tools—ex ante approval, ex post withdrawal, and liability. From the results on authorization regulation, it is clear that when $e^G < \hat{e}$ it can be useful to use liability to discourage experimentation. As we show, for $e^G < \hat{e}$, at the optimal mix of tools the approval and withdrawal standards are set at their first-best levels, while liability is used to discourage experimentation and to induce the first-best experimentation decision by the firm. When instead $e^G > \hat{e}$, it becomes pressing to encourage experimentation, liability is no longer useful, and the planner sets lenient approval and withdrawal standards.

	Approval Regulation	Withdrawal Regulation	Liability
Product safety	little	little	important
Aircraft	certification (lenient)	surveillance (some)	important
Pharmaceutical	clinical trials	surveillance (some)	rare
Business practices	limited (except mergers)	limited	antitrust litigation

Table 1: Regulatory environments across applications.

Section 4 relates these results to the preemption doctrine, which contends that regulatory approval should shield firms from future liability claims. In *Wyeth v. Levine*, 555 U.S. 555 (2009), the US Supreme Court held in a 6-3 vote that FDA approval for marketing and labeling of a medication does not shield the manufacturer from product liability lawsuits under state law. Proponents of the preemption doctrine argued that exposing pharmaceutical companies to liability would reduce innovation incentives prior to market introduction. Opponents maintained that the threat of liability would induce more timely information disclosure and voluntary withdrawals of harmful drugs; see Garber (2013). For medical devices, on the contrary, the preemption doctrine prevailed following the 8-1 decision in *Riegel v. Medtronic*, 552 U.S. 312 (2008). When striking the balance between regulation and compensation for preemption cases ranging from drugs to road safety, the Supreme Court tends to side with the views expressed by the regulatory agencies, as argued by Sharkey (2008). This way, the solution can be fine tuned to the specific features and parameters of different sectors.

Our analysis justifies the preemption doctrine only when the positive externality in the good state is large. We show that optimal regulation is lenient: the approval and withdrawal standards are always set below their first-best levels. Regulation deviates from the first best only when liability should no longer be used, i.e., when experimentation should be encouraged, in which case the planner is induced to be lenient. This prediction correspond to patterns observed across a range of applications, as summarized in Table 1. Regulatory agencies tend to be lenient whenever they are involved in the process. For instance, in the case of aircraft regulation, since 2005 tests and certification for a large number of elements are delegated to the industry under the Organizational

Designation Authorization (ODA) program (Jakubiak 1997). Even for drugs, where ex ante testing is quite strict, ex post surveillance is difficult to implement (Berniker 2001 and Han et al. 2017), preventing the agency from ordering timely withdrawal. The first three rows of Table 1 rank our applications to safety regulation in terms of the level of the positive externalities: absent for most products, intermediate for aircraft that improve transportation, and large for drugs with sizeable health benefits. As the positive externality e^G increases, the mix becomes more reliant on authorization regulation relative to liability.

Our results also shed light on the regulation of new drugs. With a new disease such as COVID-19, the externality is large, once the reduction in the direct damage to public health is compounded with the indirect economic benefit as social distancing measures are phased out. According to our results, approval should occur earlier, moving quickly to ex post learning, and liability should be lifted. As a matter of fact, pressure has been put on the FDA and EMA to accelerate approval. In addition, a number of governments have committed to removing liability and granting vaccine companies legal indemnity. The crisis also suggests the phases of experimentation should be more flexibly tailored to the parameters (such the size of the externality in the good state) prevailing in different situations.

Contribution to Literature. The decision-theoretic social benchmark analyzed in Section 2 builds on Wald's (1945) sequential information acquisition model, gaining analytical traction by adopting the continuous-time formulation pioneered by Dvoretsky, Kiefer, and Wolfowitz (1953). While the Wald model entails a single phase with the adoption of an *irreversible* action, we innovate by adding a second phase in which learning takes place at a possibly different rate and adoption can be withdrawn.[6] This implies that payoffs from approval are not fixed but state dependent, thus generalizing the classic Wald problem. Section 2.1 isolates this first contribution by presenting the baseline decision-theoretic version (corresponding to the first-best benchmark) of our two-stage model with experimentation and learning. Learning in the ex post phase takes place only if the safe arm represented by the withdrawal decision is not pulled, as in the bandit

[6] While Moscarini and Smith (2001) extend Wald's baseline model to allow the decision maker to undertake multiple experiments in each instant, we instead allow for different experimentation stages, but still with a single experiment in each instant. Within a setting with a single irreversible decision, Che and Mierendorff (2019) push Wald's framework by allowing the decision maker to choose signal structures that favor learning about a state over the other. Zhong (2019) endogenizes all aspects of learning, beyond precision and direction.

literature.[7,8]

This decision-theoretic benchmark is also relevant to model consumer search for product information. The ex ante phase corresponds to Branco, Sun, and Villas-Boas (2012): the potential buyer chooses between searching for more information, abandoning search, or purchasing the product. In addition, our framework allows the consumer to keep learning about product attributes after purchase, with the option of either returning the product or reselling on the secondary market following bad information.[9]

Our second contribution is the analysis of the strategic problem of dynamic regulation of an activity with uncertain externality. The case with irreversible approval is related to Carpenter's (2004) pioneering application of experimentation models to approval regulation. More directly, we build on multi-agent extensions of the Wald model proposed by Chan, Lizzeri, Suen, and Yariv (2018) and Henry and Ottaviani (2019), to which we add ex post learning and the possibility of reverting the decision.[10] By modeling information management in regulation as a dynamic persuasion problem, we also link to Kamenica and Gentzkow (2011) and the recent follow-up literature. In this vein, Orlov, Skrzypacz, and Zryumov (2020) analyze dynamic information control by an agent who aims at persuading a principal to delay withdrawal—in our setup we can analyze information arrival both before and after adoption up to withdrawal.

In terms of applications, our normative analysis of optimal regulation connects us to the law and economics literature on safety regulation and liability. The bulk of this literature focuses on incentivizing firms to optimally invest in ex ante precautions so as to limit the negative externalities generated by risky activities.[11] Our dynamic setting allows for both ex ante precautions

[7] Bolton and Harris (1999) introduced strategic issues into bandit models; closer in spirit to us Strulovici (2010) analyzed a model of collective experimentation in which decision makers jointly control actions that result in information. For recent analyses of agency models of experimentation see also Green and Taylor (2016), Guo (2016), Halac, Kartik, and Liu (2016), and Grenadier, Malenko, and Malenko (2016).

[8] For the baseline canonical setting with irreversible decision, Morris and Strack (2017) show that any posterior-separable cost of information can be written as the outcome of a sequential sampling model, as in Wald's setup. Pomatto, Strack, and Tamuz (2019) axiomatically characterize all cost functions with constant marginal cost of information.

[9] The model could be extended to include the choice of product price (as in Branco, Sun, and Villas-Boas 2012) as well as a buy-back price; see also Inderst and Ottaviani (2013).

[10] In the context of the baseline setting with irreversible decision, McClellan (2017) characterizes the commitment solution, while Bizzotto, Rudiger, and Vigier (2020 and 2021) allow for the arrival of outside information and for the receiver to be privately informed.

[11] For example, Shavell (1984) and Kolstad, Ulen, and Johnson (1990) show that a mix of liability and ex ante regulation is welfare improving whenever injurers can escape suit or court's behavior is uncertain. In the same vein, Schmitz (2000) shows that a mix of liability and ex ante regulation is optimal if wealth varies among injurers. Friehe and Schulte (2017) consider the optimal mix between liability and ex ante regulation in a setting where the firm can run a specific experiment before asking for approval. For broader overviews of the economic analysis of product liability see Polinsky and Shavell (2010) and Daughety and Reinganum (2013).

(corresponding to the approval standard) and ex post precautions (the withdrawal standard), as well as for experimentation.[12] Ex ante precautions impact the experimentation standard. In turn, ex post precautions affect the ex ante precautions taken. In law and economics, Schwartzstein and Shleifer (2013) also emphasize the dynamic distortions arising from the gap between social and private returns to economic activity; while in their model firms are privately informed about their type (safe or risky), in our model firms experiment and learn alongside the regulator. A key role played by liability in our setting is to discourage socially excessive experimentation. Our results also speak to the empirical literature studying how liability chills innovation. Using an increase in liability risk in the medical implant market, Galasso and Luo (2018) establish an adverse effect on upstream innovation, in contrast with some earlier findings by Viscusi and Moore (1993) and Galasso and Luo (2017).

Oi (1973) and Hamada (1976) are among the first to notice that liability can impact the willingness to pay of consumers—our model abstracts away from this effect.[13] Hua and Spier (2020) show that firms with market power find it optimal to under-provide safety and to disclaim responsibility for consumer harm, provided that consumers that use the product more intensely are also more likely to suffer harm. Their result offers a rationale for liability regulation, even when firms set prices and are allowed to compensate consumers for the damages generated by the products they sell. Taking as a given the need for regulation, instead, we compare different forms of intervention and characterize the optimal regulatory mix of instruments.

Viscusi, Vernon, and Harrington (1995, pages 785-786) informally discuss the ex ante and ex post modes of government regulation and point out how liability may chill research incentives. They highlight that "a final issue on the policy agenda is the overall coordination of regulatory and liability efforts." This is exactly the interaction we characterize.

[12] Ottaviani and Wickelgren (2009) and (2011) offer a complementary modeling approach based on signaling of the firm's private information in the context of a two-period model. In the area of competition policy, Rey (2003, Section 4.2) informally discusses the pros and cons of ex ante regulation v. ex post antitrust.

[13] The literature on product recalls also focuses on contracting between the firm and consumers; see e.g. Welling (1991), Marino (1997), and Rupp and Taylor (2002). In particular, Spier (2011) shows that, even under strict liability—whereby any harm caused is fully compensated by the firm—the buyback price is inefficiently low, because of the firm's monopsonistic position in the ex post stage when products are recalled.

2 Balancing Ex Ante Experimentation and Ex Post Learning

2.1 Planner Benchmark: Model

In our welfare benchmark, the social planner p decides whether to introduce a product under uncertainty about the state of the world θ, which can be either good G or bad B.[14] The prior belief about the state is $q = \Pr\{\theta = G\}$. Information arrives in two phases. In the first phase, at each instant a decision needs to be made: either reject R, experiment E, or adopt A. Rejection is irreversible: the game ends following R and the planner obtains a zero payoff. Experimentation costs c per unit of time and results in the public revelation of information, as explained below. Adoption ends the first, ex ante phase of the game and starts the second, ex post phase in which at each instant of time the planner needs to decide whether to continue C (being active in the market) or withdraw W (ending the game).[15] While active in the market, following adoption and up until withdrawal, the social planner collects state-dependent flow payoff $v_p^G > 0$ in the good state and $v_p^B < 0$ in the bad state. Withdrawal is irreversible and gives the planner a payoff equal to zero. The planner discounts future payoffs at rate $r \geq 0$.

Information Arrival. We now describe information arrival in the two phases, $n \in \{\mathrm{I}, \mathrm{II}\}$, where $n = \mathrm{I}$ denotes the ex ante experimentation phase when E is played and $n = \mathrm{II}$ the ex post learning phase when C is played. The arrival of new information is modeled as a Wiener process $d\Sigma$ with variance ρ^2 and state-dependent drift: positive drift μ_n in state G and negative drift $-\mu_n$ in state B. The ex ante and ex post phase differ, not only in the speed of accumulation of information μ_n, but also in the cost of research.

While information collection over a period of time dt entails costs cdt in the ex ante stage, it only has an indirect and state-dependent cost in the ex post stage based on the adoption payoff, given that the planner's payoff in the bad state is $v_p^B < 0$. The state-dependent flow payoffs v_p^G and v_p^B are not observed. Only the realizations of the stochastic process is observed at time $t > 0$ and it is used to update the belief. The updated belief is a sufficient statistic for all the information collected up and until that instant of time.[16]

[14]Fudenberg, Strack, and Strzalecki (2018) make strides in the analysis of a Wald experimentation problem with a richer state space, but still with only one (ex ante) phase.

[15]The model can be easily extended by making the withdrawal option costly.

[16]To simplify derivations we express beliefs in terms of the log-odds ratio $\sigma_t = \log \frac{q_t}{1-q_t}$. As explained in Online Appendix A, the log odds posterior belief follows a Brownian motion.

2.2 Planner Solution

We now characterize the first-best solution of the planner problem. This analysis serves as a welfare benchmark as well as a building block for the rest of the paper.

Ex Post Learning: Bandit Problem. In the ex post phase, the planner faces a bandit problem.[17] At each instant t, the choice is between a safe arm (withdrawal W with payoff 0) and a risky arm (continue C with expected payoff equal to $q_t v_p^G + (1-q_t)v_p^B$, where q_t is the posterior belief that $\theta = G$ conditional on the information available at time t).

When the current belief is q, the value for the planner in the ex post phase is equal to the payoff the planner expects to accumulate before potentially withdrawing if the belief reaches z,

$$u_p^{\mathrm{II}}(q) = q u_p^{\mathrm{II}}(q,G,z) + (1-q) u_p^{\mathrm{II}}(q,B,z), \tag{1}$$

with ex post conditional values given by

$$u_p^{\mathrm{II}}(q,\theta,z) = \frac{v_p^\theta}{r}\left(1 - \psi^{\mathrm{II}}(q,\theta,z)\right)$$

where $\psi^{\mathrm{II}}(q,\theta,z) \equiv E[e^{-rT_z}|\theta,q]$ is the expected discounted value of receiving a payoff of 1 when the belief hits for the first time the withdrawal standard z in state θ starting from belief q. T_z is the first time the belief hits threshold z.[18] The ex post conditional value $u_p^{\mathrm{II}}(q,\theta,z)$ corresponds to the flow payoff v_p^θ collected before withdrawal. This payoff is smaller in state B than G for two reasons: the flow payoff v_p^θ is smaller and the withdrawal standard is reached more quickly i.e., $\psi^{\mathrm{II}}(q,G,z) < \psi^{\mathrm{II}}(q,B,z)$.

As Proposition 1 shows, there is a unique withdrawal standard z^* that solves

$$q\frac{\partial u_p^{\mathrm{II}}(q,G,z)}{\partial z} + (1-q)\frac{\partial u_p^{\mathrm{II}}(q,B,z)}{\partial z} = 0. \tag{2}$$

This withdrawal standard is independent of the current belief. The top panel of Figure 1 represents the ex post value function as the dashed-dotted red curve. The planner keeps staying on the market for beliefs above z^* and withdraws as soon as the belief falls below that standard. At z^* the smooth-pasting condition applies, so that the ex post value function is tangent to 0 at z^*.

[17] The ex post problem is analogous to the classic continuous-time two-armed bandit problem (see, for example, Berry and Fristedt 1985 and Bolton and Harris 1999) which is also equivalent to a one-time one-option optimal stopping problem.

[18] Online Appendix A reports closed-form expressions for $\psi^{\mathrm{II}}(q,\theta,z)$.

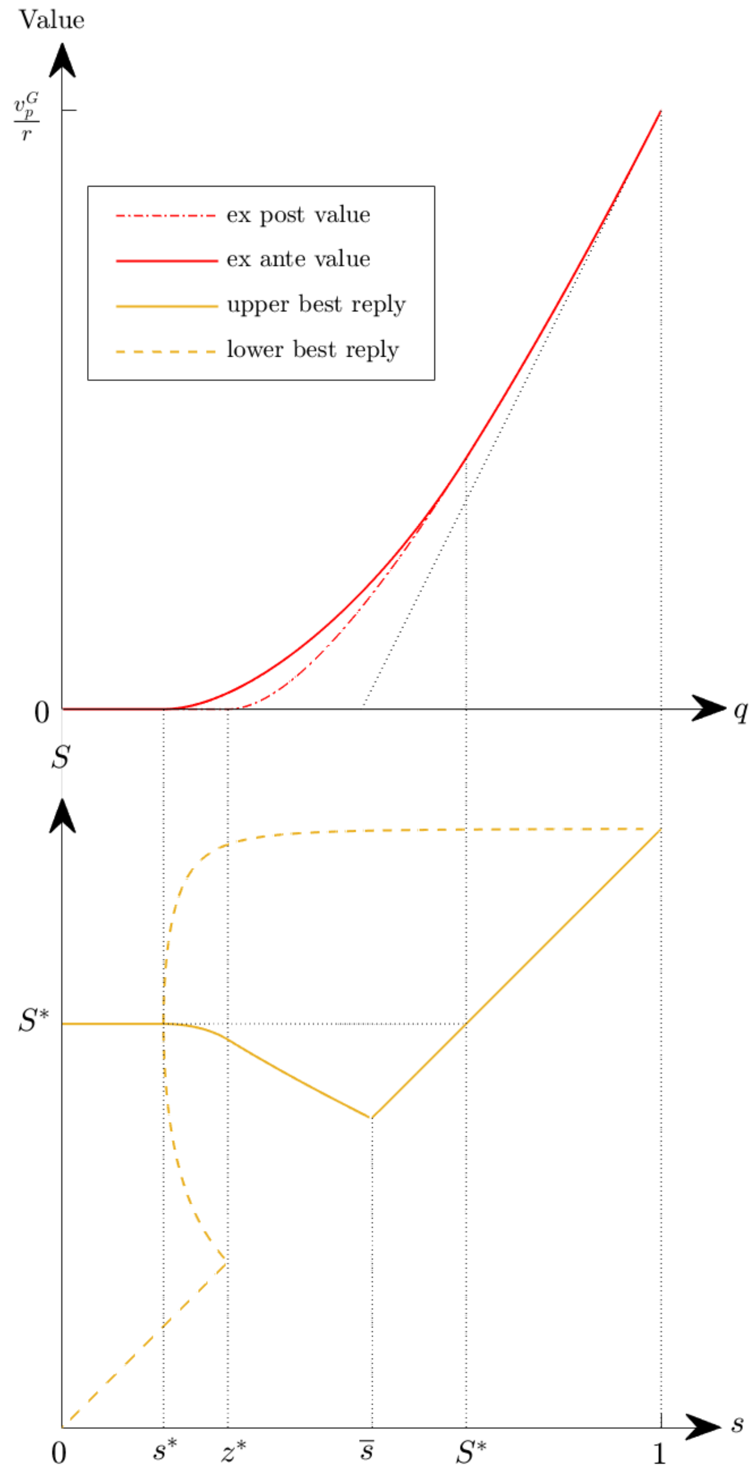

Figure 1: Top panel: ex ante value function in solid red, ex post value function in dashed-dotted red. Bottom panel: upper best reply in solid dark yellow and lower best reply in dashed dark yellow.

Ex Ante Experimentation: Reversible Wald. From the perspective of the ex ante phase, the optimal withdrawal z^* pins down the expected payoff of the planner upon adoption. In the ex ante phase, the planner chooses a lower standard s, corresponding to the posterior belief at which the planner rejects, and an upper standard S, the posterior belief at which the planner adopts.

Given a set of standards (s,S), we now derive the expected payoff of the planner. Denote 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 $\Psi^{\mathrm{I}}(q) \equiv E[e^{-rT_S}|q]$ as the expected discounted value in the ex ante phase of receiving a payoff of 1 when the belief hits for the first time standard S, conditional on not hitting standard s before. Similarly, the expected discounted probability $\psi^{\mathrm{I}}(q) \equiv E[e^{-rT_s}|q]$ is the expected discounted value of a payoff of 1 received when the belief hits for the first time s, conditional on not hitting S before.

The planner obtains the second-stage expected payoff of $u_p^{\mathrm{II}}(S)$ if S is reached before s; starting from q this happens with probability $\Psi^{\mathrm{I}}(q)$. With probability $\psi^{\mathrm{I}}(q)$, the second-stage payoff is $u_p^{\mathrm{II}}(s)$. During the ex ante experimentation phase, the planner pays an instantaneous research cost c, and thus c/r in discounted present value, until the belief hits either s or S, which happens with probability $1-\psi^{\mathrm{I}}(q)-\Psi^{\mathrm{I}}(q)$. Overall, the ex ante expected payoff of the planner is

$$u_p^{\mathrm{I}}(q) = \psi^{\mathrm{I}}(q)u_p^{\mathrm{II}}(s) + \Psi^{\mathrm{I}}(q)u_p^{\mathrm{II}}(S) - \left[1-\psi^{\mathrm{I}}(q)-\Psi^{\mathrm{I}}(q)\right]\frac{c}{r}. \tag{3}$$

At the solution of the problem, the planner experiments for intermediate beliefs ($s^* < q < S^*$) and either rejects following the accumulation of sufficient bad news ($q \leq s^*$) or adopts following sufficient good news ($q \geq S^*$). This is not a standard Wald problem because the state contingent payoff at approval $u_p^{\mathrm{II}}(S)$ depends on the value of S.[19]

Overall Solution: Reversible Wald Followed by Bandit. The planner strikes a balance between ex ante experimentation and ex post learning. Even though information is accumulated at no direct cost in the ex post phase, there is an indirect cost due to the negative payoff in the bad state, providing room for experimentation in the ex ante phase. As shown in Figure 1, at belief z^*, the planner expects a zero payoff from adoption. Experimentation is thus valuable in the range (s^*,S^*), an interval that contains z^*.

For beliefs in (s^*,S^*), the possibility of ex ante experimentation increases the value for the planner, as shown in the top panel of Figure 1 where the ex ante value (solid red line) is above the ex post value (dashed-dotted red). The ex ante value is tangent to the ex post value exactly at S^* and to the zero horizontal line at s^*. Overall, we have:

[19] In the standard Wald problem, a given fixed value is obtained upon approval.

Proposition 1 *The first-best solution consists of three standards, for rejection s^*, adoption S^*, and withdrawal z^*, with $s^* \leq z^* \leq S^*$, such that:*

(a) *In the ex ante phase, the planner rejects if $q \leq s^*$, experiments if $s^* < q < S^*$, and adopts if $q \geq S^*$ and in the ex post phase, the planner withdraws if $q \leq z^*$;*

(b) *All standards are independent of the current belief q.*

Next, we describe the first-order conditions—which can also be interpreted as best replies—leading to the solutions. This construction plays a central role throughout the paper, not only for the planner solution but also for the analysis of the strategic interactions between the planner and the firm.

First-Order Conditions as Best Replies. To understand the incentives in the ex ante phase for a given ex post withdrawal standard z, it is useful to decompose the optimal choice of (s,S) into two steps:

- We start by characterizing the optimal choice of the rejection standard s for a given adoption standard S. In the strategic problem this constrained solution corresponds to the best reply for a player who chooses s for given S (and z) set by another player. We denote this *lower best reply* as $b_p(S)$, suppressing the dependence on the withdrawal standard z.[20]
- Similarly, the *upper best reply* is the optimal adoption standard $S = B_p(s)$ when the rejection standard is constrained to be s. We adopt the convention that the player who controls approval—and thus determines the upper best reply—can approve at any belief above the rejection standard s, and in particular can potentially approve and immediately withdraw by setting $S = B_p(s) < z$.[21]

The bottom panel of Figure 1 plots the best replies when ex post withdrawal is at the first-best level z^*.[22] The solution to the first-best problem in the ex ante phase, (s^*, S^*), lies at the intersection of the lower (dashed dark yellow) and upper (solid dark yellow) best replies, as illustrated by the top and bottom panels of Figure 1. At this intersection, $s = s^*$ is such that the ex ante value function is tangent to the zero horizontal line and $S = S^*$ is such that the ex ante and ex post values are tangent. We now provide a heuristic derivation of these two best replies.[23]

[20] In the text we focus on the case with $s < z^*$ where $u_p^{\mathrm{II}}(s) \equiv 0$; derivations for the other case are presented in Online Appendix A.

[21] In this sense, the player has control on rejection.

[22] Withdrawal z in the second phase does not depend on the choice of s or S.

[23] See Lemma 2 in Online Appendix A for a more formal analysis.

Lower Best Reply. Differentiating (3) with respect to s at $s < z^*$, when interior, the lower best reply $s = b_p(S)$ is implicitly defined by the first order condition

$$-\underbrace{\frac{\partial \Psi^{\mathrm{I}}}{\partial s}}_{<0} u_p^{\mathrm{II}}(S) = \underbrace{\left(\frac{\partial \psi^{\mathrm{I}}}{\partial s} + \frac{\partial \Psi^{\mathrm{I}}}{\partial s}\right)}_{>0} \frac{c}{r}. \tag{4}$$

The player controlling experimentation equalizes the marginal benefit of experimentation on the left-hand side, which increases the probability of reaching approval and obtaining $u_p^{\mathrm{II}}(S)$, with the marginal cost of experimentation on the right-hand side. As displayed in Figure 1, for $S \leq z^*$, there is no value of experimentation, since reaching S yields a zero payoff. Thus, for $S \leq z^*$, there is no experimentation and the lower best reply follows the diagonal ($s = S$). For $S > z^*$, the planner starts experimenting (i.e. $s < S$), trading off the direct experimentation cost with the value of information generated as described in equation (4). Initially, the value of experimentation quickly increases with S, so that the lower best reply decreases with S down to a minimum s^* (for $S = S^*$). In this case adoption and experimentation are strategic substitutes.[24] For higher value of S, the best reply $s = b_i(S)$ increases with S to reach 1 in the limit as $S \to 1$. In this case the main driving force of the decision to experiment is the expected financial cost of reaching S, which increases with S. Adoption and experimentation are then strategic complements.

Upper Best Reply. Differentiating (3) with respect to S, when interior the upper best reply $S = B_p(s)$ is implicitly defined by the first order condition

$$\Psi^{\mathrm{I}} \frac{\partial u_p^{\mathrm{II}}(S)}{\partial S} = -\underbrace{\frac{\partial \Psi^{\mathrm{I}}}{\partial S}}_{<0} u_p^{\mathrm{II}}(S) - \underbrace{\left(\frac{\partial \Psi^{\mathrm{I}}}{\partial S} + \frac{\partial \psi^{\mathrm{I}}}{\partial S}\right)}_{<0} \frac{c}{r}. \tag{5}$$

The player controlling adoption equalizes the marginal benefit of delaying adoption on the left-hand side with the marginal cost on the right-hand side. The former is positive because, upon approval, it is more likely that the state is good when S is set higher. The latter combines the opportunity cost of delaying approval with expected payoff $u_p^{\mathrm{II}}(S)$ as well as the financial cost of experimentation. Under the convention we imposed, this player can always ensure himself the ex post value (red dashed-dotted) by approving and immediately withdrawing at s^*, approving at S^*, and experimenting in (s^*, S^*).

The bottom panel of Figure 1 plots the upper best reply in solid dark-yellow, with the initial belief q_0 set at $q_0 = z^*$. For $s \in [0, s^*]$, $B_p(s)$ is constant at level S^*.[25] Indeed for $s < s^*$, if from

[24] Increasing S reduces adoption rates while reducing s increases experimentation by reducing rejection rates.

[25] In Figure 1 for $s \leq s^*$ we chose the convention to plot $B_p(s) = S^*$. Alternatively, we could have chosen to plot $B_p(s) = s^*$.

q_0 bad news is collected and belief q falls, it is too costly to experiment until s and preferable to adopt (and immediately withdraw) at s^*, thus interrupting experimentation. If, instead, good news is accumulated and q increases, it is optimal to approve at S^*.

For $s \geq s^*$ the upper best reply is interior and balances the marginal value of information with the marginal costs of experimentation and delayed adoption, according to equation (5). As s increases from s^*, the upper best reply $B_p(s)$ decreases toward the diagonal. In this case, the experimentation cost plus opportunity cost of delaying approval are the main concern—adoption and experimentation are strategic complements. Finally, for $s \in (\bar{s}, 1]$ there is no non-empty interval of beliefs for which the ex ante value would be above the ex post value—the upper best reply lies on the diagonal $B_p(s)$.

3 Regulating Private Activity

3.1 Introducing the Firm: Model

We now enrich the model by introducing a firm f that interacts strategically with the planner p. In the ex ante phase, the firm bears the full experimentation cost cdt for every instant dt of active experimentation. The information collected is publicly revealed.[26] In the ex post phase, the firm collects a flow payoff v_f, with $0 \leq v_f \leq v_p^G$, for every instant of active presence in the market after adoption and until withdrawal. We take v_f as exogenous; in particular, v_f is not affected directly by the level of liability or other regulatory tools.[27]

To simplify derivations, the firm's payoff v_f is assumed to be independent of the state. This assumption is sensible in situations in which consumers, in spite of the losses they bear in the bad state, are not able to react to the information that arrives over time. The setting captures environments in which information is technical and hard to process for consumers. It is also reasonable in contexts where the adverse effects in the bad state happen in bulk, an extension we consider in Online Appendix D. Making the firm's baseline flow payoff v_f state dependent would affect the results by providing a channel, other than regulation, to discipline the firm.[28]

[26] In the case of research on drugs, the public information assumption does not appear unreasonable since before market introduction, firms have to reveal the results of clinical trials and are legally obliged after introduction to provide further information they obtain on side effects. The model can be extended to allow for private information collection with costly lying as in Kartik, Ottaviani, Squintani (2007), where the expected cost depends on the size of the lie.

[27] If, instead, the interaction between the firm and consumers were to be modeled explicitly, a change in liability or in other regulatory tools would directly affect the willingness to pay of consumers and thus v_f. Future work could investigate the robustness of our insights when the payoff of the firm is obtained endogenously from trading with consumers, as in Hua and Spier (2020).

[28] Our results would still hold as long as the gap in payoff between states is smaller for the firm than for the planner:

Absent regulation, incentives of firm and planner are misaligned because the firm does not suffer all the negative social cost in the bad state, but also cannot recover the full social benefits in the good state.[29] The firm's activity generates a positive externality $e^G = v_p^G - v_f \geq 0$ in the good state and a negative externality $e^B = v_p^B - v_f < 0$ in the bad state on the rest of society.[30]

Left to its own devices under laissez-faire, the firm would not experiment, adopt immediately, and never withdraw. To align the incentives of firms with societal goals in areas ranging from safety regulation to competition policy, over the years legislators introduced a number of instruments. The initial approach has been to use ex post instruments such as liability following harm and, then, ex post regulation with mandatory withdrawal of dangerous products. With the rise of the regulatory state from the beginning of the 20th century, ex ante regulation has been introduced in many areas by instituting a structured ex ante authorization process and delegating its implementation to specialized regulatory agencies, as described by Glaeser and Shleifer (2003).

Sections 3.2, 3.3, and 3.4 analyze how regulation through liability, withdrawal, and authorization affect the margins for experimentation z, adoption S, and withdrawal z.

3.2 Liability

Historically, liability was first introduced to compensate victims following the occurrence of an adverse event, but also to discipline firms and influence their ex ante choices. We model this instrument in a reduced-form way by allowing the planner to commit at date 0 to a liability rate $L \geq 0$ imposed on the firm per unit of time on the market if $\theta = B$.[31] In state B, the firm thus obtains payoff $v_f - L$; in state G, instead, the payoff remains unaffected at v_f. Liability is strict in the sense that it does not depend on actual negligence or intent to harm by the firm. The assumption that the liability rate L is imposed on the firm per unit of time on the market captures the idea that damages paid in practice are often calculated in proportion to the harm caused, and in the bad state, the longer the product is on the market, the more it is the case.[32] We denote the

$v_f^G \leq v_p^G$ and $v_f^B \geq v_p^B$ with $v_f^B \geq 0$.

[29]Think of positive externalities in consumption, as in the case of vaccines. Misalignment can also arise from consumer heterogeneity, as in Hua and Spier (2020). Future work could extend the analysis to endogenize v_f as arising from the interaction between the firm and consumers.

[30]While we focus on comparative statics with respect to e^G, our results can be obtained from comparative statics with respect to $v_f \in [0, v_p^G]$. For instance, it is natural to expect that stronger (weaker) market competition would reduce (increase) v_f and thus increase (reduce) e^G.

[31]We purposefully abstract away from the details of the judicial procedure. Given risk neutrality, L captures the *expected* penalty the firm incurs per unit of time in the market, conditional on state is $\theta = B$. It is not necessary for the state to ever be observed—it is enough for some information about the state to be available.

[32]Many of our results still hold if news in the ex post phase are modeled as a Poisson arrival process and liability as a fixed payment following the occurrence of an adverse event, as we show in Online Appendix D. However, in this

choices when the planner commits to L by $(s_{ff}(L), S_{ff}(L), z_{ff}(L))$ and the optimal liability by L^*_{ff}.[33]

The planner can indirectly control the firm's ex post decision by setting the liability rate L that induces any withdrawal standard z. In particular, $\hat{L}$ denotes the liability rate that induces first-best withdrawal, $z_{ff}(\hat{L}) = z^*$. The choice of liability also indirectly affects the two ex ante decisions, experimentation and adoption. The optimal liability trades off ex ante with ex post distortions as follows:

Proposition 2 *With strictly positive externality in the good state, $e^G > 0$, the optimal liability L^*_{ff}*

(a) *does not fully compensate for the harm caused $v_f < L^*_{ff} < -e^B$,*

(b) *is such that the firm rejects too early $s_{ff}(L^*_{ff}) > s^*$ and adopts too early $S_{ff}(L^*_{ff}) < S^*$,*

(c) *is such that there exists belief $\hat{q}$ such that the firm withdraws earlier than the first best $z_{ff}(L^*_{ff}) > z^*$ if $q > \hat{q}$ and later otherwise.*

According to Proposition 2.(a), when $e^G > 0$ the optimal liability does not require the firm to fully compensate the damage caused, which would entail choosing liability $\bar{L} = -e^B$; by setting this *full liability* level, the social planner perfectly aligns the firm's incentives in the bad state. If there is no positive externality, $e^G = 0$, the first best is achieved with full liability, $\bar{L}$. If $e^G > 0$, however, full liability no longer induces the first best. Given that the firm's payoff in the good state is socially insufficient, the firm rejects too early, $s^* < s_{ff}\left(\overline{L}\right)$, adopts too late, $S^* < S_{ff}\left(\overline{L}\right)$, and withdraws too early, $z^* < z_{ff}\left(\overline{L}\right)$. The planner should thus constrain the courts to be more lenient, as shown in result (a).

Another option is to set the liability at $L = \hat{L}$ so as to induce withdrawal at the first-best level z^*. This liability is such that the ratio of the firm's payoff relative to the social payoff is the same both in the bad and in the good state

$$\frac{v_f - \hat{L}}{v_p^B} = \frac{v_f}{v_p^G}.$$

Given that this ratio is less than 1, because the firm cannot capture the full social benefits in the good state ($v_f < v_p^G$), the firm's expected payoff is scaled down compared to the planner by a factor v_f/v_p^G. The bottom panel of Figure 2 plots the firm's best replies when $L = \hat{L}$ (i.e., when the firm withdraws at z^*) in comparison to the planner's best replies, already displayed in Figure 1.

alternative model, once the adverse event occurs the harm is already done and liability would not induce the firm to withdraw endogenously, so that the planner would need to force withdrawal.

[33]The convention we use for notation is the following: for a given standard $x \in \{s, S, z\}$, we denote by x_{ij} the standard resulting in the regulatory environment where player $i \in \{f, p\}$ controls ex ante adoption and player $j \in \{f, p\}$ controls ex post withdrawal.

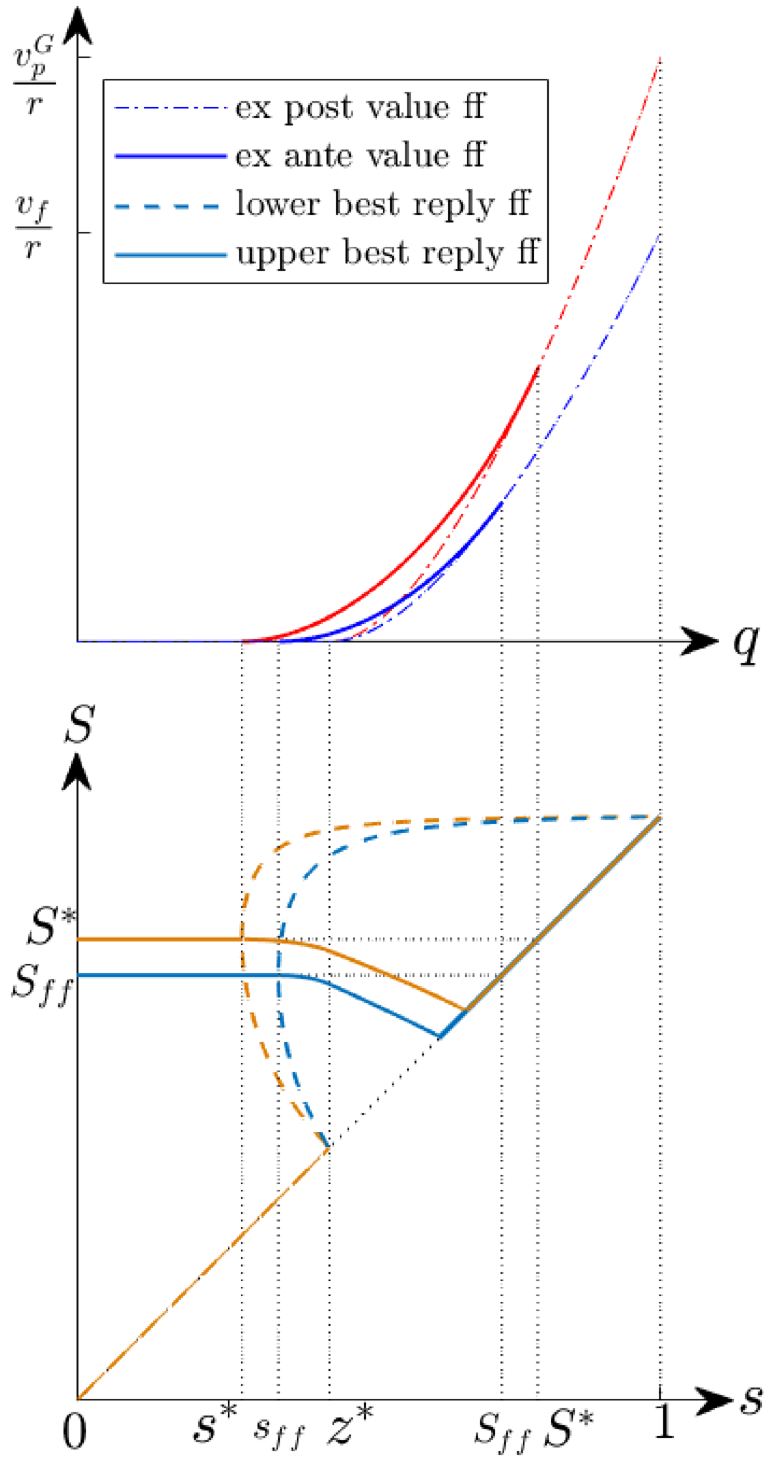

Figure 2: Value functions (top panel) and best replies (bottom panel) for the planner p and the firm f facing liability $L = \hat{L}$.

On the one hand, the firm's lower best reply $b_f(S)$ (dashed-blue) is shifted to the right compared to the planner's lower best reply $b_p(S)$ (dashed-orange). Indeed, the financial cost of experimentation is identical for both players, but the value of experimentation on the left-hand side of (4) is scaled down for the firm. Intuitively, for any given S, the firm has less incentives to experiment since the expected payoff is smaller. On the other hand, the firm's upper best reply $B_f(s)$ (solid blue) is shifted down compared to the planner's $B_p(s)$ (solid orange), since the firm attaches a lower value to information. Overall, this implies that the firm rejects too early, $s^* < s_{ff}(\hat{L})$, and adopts too early $S_{ff}(\hat{L}) < S^*$.

Even though the planner can set the liability to induce the firm to withdraw at the first-best level in the ex post phase, it is optimal for the planner to deviate from $\hat{L}$. From the ex post perspective, moving away from $\hat{L}$ induces a second-order loss, as withdrawal is moved away from the ex post optimal level z^*. Deviating from this level generates two contrasting first-order effects in the ex ante phase by affecting the incentives for experimentation s and adoption S. An increase in the liability above $\hat{L}$ induces a decrease in experimentation (the lower best reply shifts to the right, resulting in a first-order social loss) and an increase in the approval benchmark as the opportunity cost of delaying approval is reduced and the value of information raised (the upper best reply shifts upward, resulting in a first-order social gain).

According to Proposition 2.(c), the tradeoff between the effect on experimentation and the effect on adoption is resolved depending on the initial belief q. If q is low, experimentation is the most pressing issue, so it is optimal for the planner to commit to be lenient by setting a lower liability to encourage experimentation. For higher values of q, instead, increasing the approval standard becomes more important and the planner should commit to be tough by setting a higher liability. According to Proposition 2.(b), regardless of which dimension (experimentation or approval) is favored, once liability is set optimally the firm still rejects too early $s_{ff}(L^*_{ff}) > s^*$ and adopts too early $S_{ff}(L^*_{ff}) < S^*$, as also illustrated in Figure 2.

3.3 Withdrawal Regulation

Next turn to ex post regulation when firms are ordered to withdraw products after sufficiently bad news has accumulated. Denote by $(s_{fp}(z), S_{fp}(z))$ the ex ante choices of the firm when the planner commits to withdraw at z. As noticed above, an appropriate liability rate can always be chosen to achieve any withdrawal standard. For any commitment to a withdrawal standard z, denote by $\tilde{L}(z)$ the liability rate that induces the firm to withdraw at z, i.e., $z_{ff}(\tilde{L}(z)) = z$.

Liability and ex post withdrawal thus are equivalent from the ex post perspective. However, we show in this section that they induce different ex post value functions and thus shape differently ex ante incentives. The ex ante choices under withdrawal regulation at z and ex post equivalent liability $\tilde{L}(z)$ compare as follows:

Proposition 3 *For any commitment z and corresponding liability $\tilde{L}(z)$, the firm:*

(a) *experiments for any value of c under withdrawal regulation z, i.e. $s_{fp}(z) < z < S_{fp}(z)$,*

(b) *experiments more under withdrawal regulation z than under liability $\tilde{L}(z)$, i.e. $s_{fp}(z) < s_{ff}(\tilde{L}(z))$,*

(c) *adopts earlier under withdrawal regulation z than under liability $\tilde{L}(z)$, i.e. $S_{fp}(z) < S_{ff}(\tilde{L}(z))$, if $\psi^{\text{II}}(S,B,z) < \bar{\psi}$.*

To understand the mechanics of these results, compare the ex post value functions for the firm under the two regimes, for a given withdrawal standard z. As illustrated in the top panel of Figure 3, the ex post value functions (plotted as a green dotted line, for withdrawal regulation, in addition to the blue dotted line under liability, as in Figure 2) are clearly equal for $q = z$ (by construction the withdrawal policy is identical) and $q = 1$. However, liability results in a lower value than ex post regulation for $q \in (z, 1)$, given that liability taxes the firm in the bad state while the expected time on the market is the same since withdrawal occurs at z in both cases.

Moreover, by smooth pasting the value function under liability is tangent to 0 at $q = z$, given that withdrawal is optimally chosen by the firm.[34] Instead, the value function under withdrawal regulation at any $z > 0$ starts off with a kink at $q = z$ because z is not optimally chosen by the firm. This kink has important implications for the ex ante behavior of the firm under withdrawal regulation, as we explain below.

The first important consequence of this kink in the value function is that the firm always experiments in the first phase under withdrawal regulation regardless of the size of the cost c, as expressed in Proposition 3.(a). Indeed the value of introducing at z is 0, since it leads to immediate withdrawal and the kink implies that the benefit of even a marginal amount of experimentation is unbounded. This intuition applies whenever the player who experiments in one phase is not the one who optimally controls withdrawal in the next.

To shed further light on the comparison between withdrawal regulation and liability, we now consider the best replies of the firm. The bottom panel of Figure 3 plots the best replies under withdrawal regulation and under liability. The lower best reply under withdrawal regulation (dashed green) is shifted to the left compared to the lower best reply under liability (dashed blue).

[34] Following the logic presented in Section 2 for the planner problem, the ex ante experimentation and approval standards chosen by the firm are determined by smooth pasting conditions similar to (4) and (5), once we replace u_p^{II} by u_f^{II}. The ex ante value function is tangent to the horizontal line at s and to the ex post value function at S.

For a given approval standard S, the firm expects higher revenue $u_f^{\mathrm{II}}(S)$ under ex post regulation because of the saved liability in the bad state. Thus, according to Proposition 3.(b), the option value of experimentation is higher under withdrawal regulation, pushing the firm to experiment more than under liability.

The choice of adoption standard is more subtle. As explained earlier, the tradeoff characterizing the upper best reply is between the marginal cost of continuing experimentation (the information acquisition cost plus the opportunity cost of delaying the approval benefits) and the marginal benefit described in more details below,

$$\frac{\partial u_f^{\mathrm{II}}}{\partial S} = \underbrace{\left(u_f^{\mathrm{II}}(q,G,z) - u_f^{\mathrm{II}}(q,B,z)\right)}_{\text{classic Wald value of information}} + \underbrace{\left(S\frac{\partial u_f^{\mathrm{II}}(q,G,z)}{\partial S} + (1-S)\frac{\partial u_f^{\mathrm{II}}(q,B,z)}{\partial S}\right)}_{\text{postponement effect}}. \tag{6}$$

The first term captures the value of information, which increases with the gap between the expected payoff in the good and bad states, a force already present in the classic Wald model. The second term incorporates the fact that approval payoffs are conditional on S, a novel feature of our setting: when S is higher, the starting belief is further away from withdrawal z, thus the payoff in the ex post phase is collected for longer. This is what we call the *postponement effect*. If the payoffs are positive in both states of nature, as in the case of withdrawal regulation, this effect is an unambiguously positive benefit. When the firm must bear some liability for being active in the bad state, the postponement effect is dampened and may even become negative for a fixed z.

Under liability regulation when the firm (or, similarly, the planner) optimally chooses the withdrawal standard, the positive effect in the good state exactly compensates the negative effect in the bad state. To see this, note that the postponement effect is given by

$$S\frac{\partial u_f^{\mathrm{II}}(q,G,z)}{\partial S} + (1-S)\frac{\partial u_f^{\mathrm{II}}(q,B,z)}{\partial S} = -\left(S\,\frac{\partial u_f^{\mathrm{II}}(q,G,z)}{\partial z} + (1-S)\,\frac{\partial u_f^{\mathrm{II}}(q,B,z)}{\partial z}\right) = 0. \tag{7}$$

First, conditional on the state, a marginal increase in S has the opposite effect compared to a marginal increase in z. Second, the optimal z maximizes the expected value in the second phase and thus implies that the postponement effect is absent under liability (as well as in the first-best benchmark of Section 2).[35]

How do the upper best replies compare in the case of withdrawal regulation and liability? Under withdrawal regulation, payoffs are higher in the ex post phase and therefore the opportunity

[35] Specifically, we have $u_f^{\mathrm{II}}(S,\theta,z) = \frac{v_f^\theta}{r}\left(1-\psi^{\mathrm{II}}(S,\theta,z)\right)$ and since $\frac{\partial \psi^{\mathrm{II}}(S,\theta,z)}{\partial z} = -\frac{\partial \psi^{\mathrm{II}}(S,\theta,z)}{\partial S}$, this implies that under liability $S\frac{\partial u_f^{\mathrm{II}}(q,G,z)}{\partial S} + (1-S)\frac{\partial u_f^{\mathrm{II}}(q,B,z)}{\partial S} = -S\,\frac{\partial u_f^{\mathrm{II}}(q,G,z)}{\partial z} - (1-S)\,\frac{\partial u_f^{\mathrm{II}}(q,B,z)}{\partial z} = 0$.

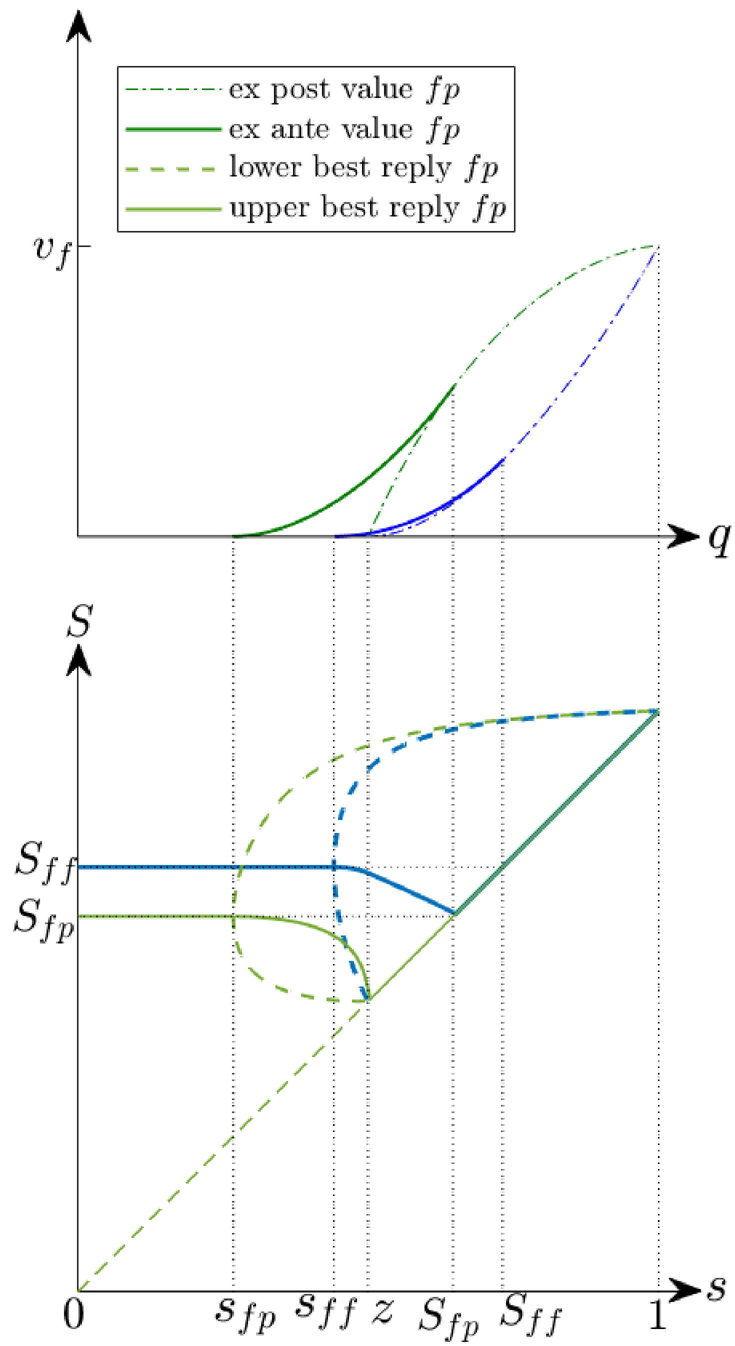

Figure 3: Value functions (top panel) and best replies (bottom panel) when the postponement effect does not dominate.

cost of delaying is larger and the value of information is lower. This would suggest that the upper best reply should be shifted downward. However, the postponement effect, absent under liability, goes in the opposite direction.

As shown in Online Appendix B where we prove result (c), the postponement effect dominates whenever ψ^{II} is sufficiently high. In this case, the opportunity cost is close to zero ($u_f^{\mathrm{II}} \simeq 0$) and so is the value of information ($u_f^{\mathrm{II}}(q,G,z) - u_f^{\mathrm{II}}(q,B,z) \simeq 0$). Thus, the only concern that remains is the postponement effect and ex post regulation leads to later approval.[36] Instead, if ψ^{II} is low, the postponement effect does not dominate and, as shown in the bottom panel of Figure 3, the upper best reply under withdrawal regulation (in dashed green) shifts downward. For the rest of the paper, we focus on this latter case and assume that ψ^{II} is low. As shown in the proof of Proposition 3, a sufficient condition is that the experimentation cost is low. This condition implies that the firm experiments a lot in the first phase, $S - z$ is large and ψ^{II} is thus small.

3.4 Authorization Regulation

In areas such as safety, environmental, occupational, and consumer protection, over the last century most Western countries have introduced some form of ex ante regulation and delegated it to regulatory agencies. Regulators have typically been granted the power to authorize and licence firms to introduce products to the market. Licences can be revoked ex post, when the agency orders withdrawal. This section studies the properties of authorization regulation, understood as the combination of ex ante and ex post regulation.

The complementarity of ex ante and ex post tools is clear. By directly targeting the ex ante approval decision, the ex post tools can instead be reoriented toward controlling withdrawal, rather than focusing on shaping ex ante incentives as in the previous sections. There is, however, one ex ante decision that neither of these tools directly controls, which is experimentation. Both ex ante regulation and ex post tools can be potentially used to discipline experimentation incentives.

We model authorization regulation as the planner committing at the outset $t = 0$ to a combination of an approval standard S_{pp} and a withdrawal standard z_{pp}. We first study the optimal authorization regulation (S^*_{pp}, z^*_{pp}) when the firm controls ex ante experimentation s_{pp}. Suppose both regulatory tools are chosen at their first-best levels, $S_{pp} = S^*$ and $z_{pp} = z^*$. For $e^G = 0$, the firm experiments too much ex ante $s_{pp} < s^*$—the firm obtains the same payoff as the planner in the good state, but does not incur the loss in the bad state and thus expects a higher payoff than the planner. When instead $e^G \to v_p^G$, the firm has no incentive to experiment and $s_{pp} > s^*$.

[36] See Figure 8 in Online Appendix B which replicates Figure 3 for the case in which the postponement effect dominates.

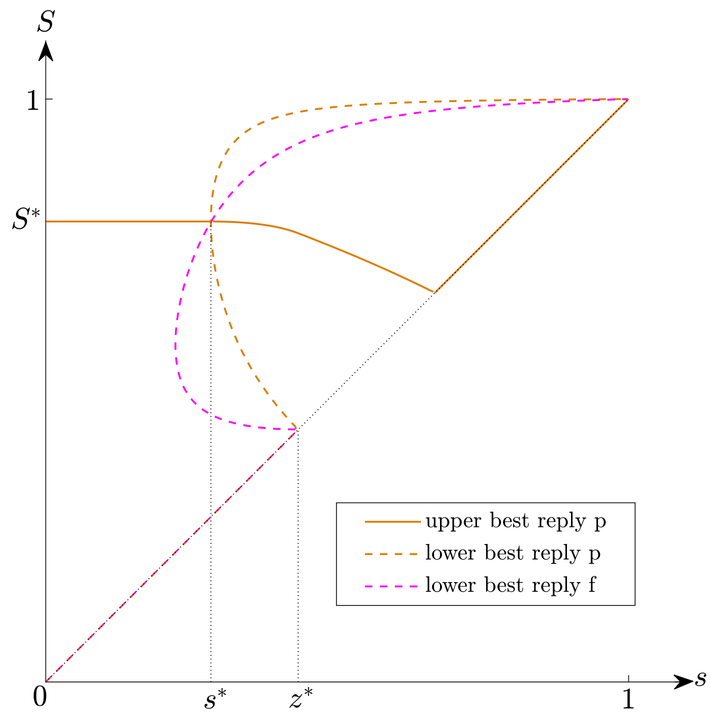

Figure 4: Best replies of planner and lower best reply of firm, for $e^G = \hat{e}$.

As e^G increases, firm f's payoff v_f decreases relatively to the social payoff v_p^G. Thus f has fewer incentives to experiment for a given value of S; the lower best reply $b_f(S)$ shifts to the left. There exists an intermediate value of e^G, denoted $\hat{e}$, such that $b_f(S)$ crosses the lower best reply of the planner precisely at S^*, as represented in Figure 4. For $e^G = \hat{e}$, the firm responds by choosing the socially optimal experimentation level, $s_{pp} = s^*$. When $e^G < \hat{e}$, the firm's lower best reply $b_f(S)$ shifts to the left; the firm's private incentives v_f are then so high that it experiments excessively relative to the first best, $s_{pp} < s^*$. If instead $e^G > \hat{e}$, incentives to experiment are insufficient, with $s_{pp} > s^*$. In both cases, the planner modifies the regulatory tools to get closer to the first best, as follows:

Proposition 4 *There exists $\hat{e}$ such that, for values of the positive externality $e^G = \hat{e}$, optimal regulation achieves the first best and:*

(a) *if $e^G < \hat{e}$, ex ante and ex post regulations are* ***stricter*** *than the first best: $z^*_{pp} > z^*$ and $S^*_{pp} > S^*$ to discourage experimentation,*

(b) *if $e^G > \hat{e}$, ex ante and ex post regulations are* ***more lenient*** *than the first best: $z^*_{pp} < z^*$ and $S^*_{pp} < S^*$ to encourage experimentation.*

For $e^G < \hat{e}$, starting from the first-best regulatory standards, the planner needs to discourage excessive experimentation. The planner achieve this by either making approval more difficult to attain or making withdrawal more likely. Proposition 4.(a) shows that in this case the planner is tougher with both standards. In both cases, moving away from the first-best levels results in second-order losses that are trumped by first-order gains in terms of experimentation incentives. According to Proposition 4.(b), if instead $e^G > \hat{e}$, the planner encourages experimentation by setting both standards at more lenient levels.

3.5 Optimal Mix of Policy Tools

How does the welfare performance of liability and authorization regulation compare? When does adding liability to authorization regulation strictly improve welfare?

Proposition 5 **(a)** ***Externality Criterion:*** *When comparing liability and authorization regulation, there exists $\tilde{e}$ such that liability welfare dominates authorization regulation if and only if $e^G \leq \tilde{e}$.*

(b) ***Preemption:*** *When the planner has access to all tools, the socially optimal mix (S_m^*, z_m^*, L_m^*) is such that liability is preempted ($L_m = 0$) if and only if $e^G \geq \hat{e}$.*

(c) ***Lenient Regulation Property:*** *For $e^G > \hat{e}$, the ex ante approval and ex post withdrawal standards at the optimal authorization regulation are more lenient than at the first best, $z_{pp}^* < z^*$ and $S_{pp}^* < S^*$. For $e^G \leq \hat{e}$, ex ante approval and ex post withdrawal standards are at the first-best levels, z^* and S^*, whereas liability is not preempted ($L_m^* > 0$) and is decreasing in e^G.*

According to Proposition 5.(a), authorization is preferred to liability for activities that generate sufficiently high externalities. Figure 5 compares the welfare performance of authorization and liability: for $e^G \leq \tilde{e}$ welfare under liability (blue line) is higher than welfare under authorization (green line), while for $e^G > \tilde{e}$ welfare under authorization dominates. Intuitively, for high e^G the planner is mainly concerned with preserving the firm's incentives to experiment and thus prefers authorization to liability.

Proposition 5.(b) characterizes the optimal regulatory structure for a planner who has access to all tools (S, z, and L). When the externality is sufficiently high, $e^G > \hat{e}$, it is welfare optimal to preempt liability. As illustrated in Figure 5, in this case the welfare obtained by combining authorization and liability (black line) is the same as the welfare under authorization alone (green line). In addition, both the approval and withdrawal standards should be more lenient than at the first best, as described in Proposition 5.(c). Intuitively, by Proposition 4 in this case experimentation is insufficient even in the absence of liability—adding liability would only reduce it further.

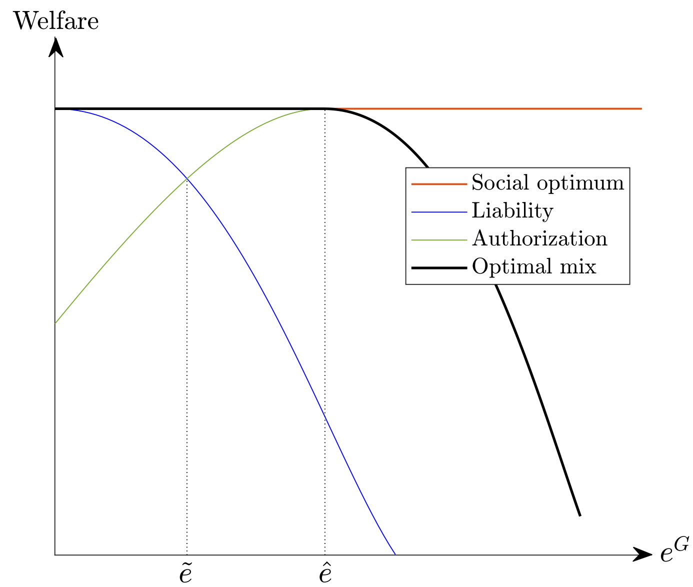

Figure 5: Welfare under different regulatory tools.

According to Proposition 5.(c), liability becomes a crucial tool for regulating activities with low externalities. When $e^G = 0$, liability alone is enough to achieve the first best. For $e^G \in (0, \hat{e})$ the first best is not attainable by either authorization or liability in isolation, but it becomes attainable once the tools are combined. As we know from Proposition 4, if the externality level is low, $e^G \leq \hat{e}$, experimentation incentives are excessive. When the planner does not have access to the lever of liability, authorization is set at stricter standards than first best. When liability is available, instead, the planner can reorient approval and withdrawal to their first-best levels and use liability to curb excessive experimentation.

4 Applications

This section presents in more detail the applications outlined in Table 1. The observed patterns are interpreted in the light of the predictions of our model.

4.1 Lenient Regulation Property

Our model applies directly to externalities which do not affect the willingness to pay of consumers, as in the case of health and safety.

Aircraft Safety. Aircraft regulation entails a mix of liability and mostly ex post regulation. Arguably, transportation generates an intermediate level of external benefits to the rest of society, taking also into account the negative impact on the environment. Even though US courts have been reluctant to examine cases for crashes that happened outside the US, in the case of the 737 MAX, Boeing faces the prospect of substantial payouts to the families of passengers if it is found responsible for both the Indonesia and Ethiopia crashes. Legal experts say the second suit could prove even more damaging, as the manufacturer was warned by the first accident (see Konert 2019 for a discussion).[37]

While for most products authorization regulation is absent, to be able to sell an aircraft the manufacturer must first be certified by the FAA. However, we argue that this regulation tends to be lenient. Since it was created, the FAA has been delegating the certification of some components to the industry. This process was accelerated in 2005 with the adoption of the Organizational Designation Authorization Program. According an Audit Report by the Office of Inspector General (2015), "one aircraft manufacturer approved 90% of the design decisions for all of its own aircraft." As a result, the certification period is relatively short. For instance, the time from first flight to certification was 394 days for the 737 MAX. Ex post withdrawal regulation is also present, but the FAA rarely uses its power to ground planes. For instance, in 1979 the FAA grounded the McDonnell Douglas DC-10. The decision was made following an American Airlines accident in Chicago with more than 270 casualties. An investigation into the accident found it was caused by maintenance issues, rather than structural problems, so DC-10 were allowed to fly again. More recently the 737 MAX was grounded, a decision directly made by the US government. In both these cases, there was no voluntarily withdrawal by the manufacturer or the airline companies.

Drugs. The authorization of new drugs is more time consuming relative to airplanes. Prior to approval, the FDA reviews the evidence contained in a battery of clinical trials conducted in a sequence of phases.[38] The FDA and the sponsor must also agree on the labelling of the product. While this ex ante regulation is lengthy and stringent, ex post regulation remains rather lenient. Post approval, the FDA requires pharmaceutical firms to report adverse events experienced by patients as well as results of post marketing clinical trials. Patients and doctors can also directly report to the FDA through the MedWatch system. The FDA at any point can request a change in the labeling of the drug or order a recall.[39] However, in practice ex post supervision is fraught

[37] Boeing has already announced that it will pay $100 million to the families, but current estimates of litigation costs are more in the order of 1 billion, in particular if airline companies need to be compensated.

[38] In terms of the model, the key state variable that determines market introduction is the efficacy of the drug relative to the risk of serious side effects. See Adda, Decker, and Ottaviani (2020) for a recent empirical analysis of how regulation shapes the incentives of the pharmaceutical industry to perform and report clinical trials across different phases.

[39] It is hard to establish empirically what proportion was voluntarily recalled by the firm and what proportion of the recalls were triggered by FDA action. In both cases, the FDA plays a key role in influencing the event through

with reporting issues and the inference from observational data is notoriously unreliable (Berniker 2001 and Han et al. 2017).[40]

Liability is rare, unless the firm is shown to have withheld information. Product liability law imposes legal obligations to compensate users injured by a defective product. The defect most commonly invoked for pharmaceuticals is a failure to warn, whereby manufacturers are liable for not issuing a warning about risks they knew, or should have known, about.[41] In cases such as Vioxx that resulted in large damages, companies failed to report to the FDA or to the medical community some evidence of adverse effects they gained after obtaining marketing approval.[42] Damage claims against pharmaceutical companies tend to be more common in the US than in other countries. According to Lybecker and Watkins (2015), for example, "there is no evidence of any court cases with positive settlement payments documented within the UK."[43]

The *Lenient Regulation Property* (authorization regulation, whenever it is used, tends to be lenient) is broadly consistent with these applications.[44] Of course, if the regulator could also use subsidies to encourage experimentation, the need for lenient regulation would be weakened. In an industry like pharma where research subsidies are common, it is then natural to expect more stringent authorization regulation.

4.2 Preemption Doctrine

The US Supreme Court has been called to rule on preemption in a number of areas, ranging from drugs (*Wyeth v. Levine*) and medical devices (*Riegel v. Medtronic*) to tobacco (*Cipollone v. Liggett Group*), road safety (*Geier v. American Honda Motor Co.*), and boat safety (*Sprietsma v.*

its surveillance program. See Zuckerman, Brown, and Nissen (2011) on the prevalence of withdrawals of drugs and medical devices.

[40]For recent empirical advances on postmarket surveillance see also Ahuja et al. (2021).

[41]In addition to the failure to warn, product liability typically addresses two other categories of defects. First, manufacturing defects, when the product does not meet the manufacturer's own design specification. Second, design defects, whereby manufacturers are liable of the foreseeable injury risks involved in the product that could have been avoided using a reasonable alternative design; see Henderson and Twerski's (1998) account of the American Law Institute Restatement (Third) of Torts. These two categories are very rarely invoked for drugs.

[42]In the Vioxx case it was judged that the warnings to physicians were inadequate in light of cardiovascular risks that became known to Merck.

[43]The essential difference is that a UK drug manufacturer can avoid litigation by arguing that it was unaware of the side effect. Lybecker and Watkins (2015) report that UK manufacturers can use the defense that "the state of scientific and technical knowledge was not such that a producer of products of the same description as the product in question might be expected to have discovered the defect if it had existed in his products while under his control."

[44]In addition to the case of aircraft and drugs, in the area of patents, the United States Patent and Trademark Office is generous both in terms of initial approvals (Jaffe and Lerner 2007) as well as withdrawals through patents invalidation, because courts apply a presumption of validity that puts a high burden of proof on the challenger (see Hall et al. 2004).

Mercury Marine). According to Sharkey (2008), roughly half of the rulings have been in favor of the preemption doctrine, suggesting that the question is far from settled.

According to Proposition 5.(b), the preemption doctrine is justified when the positive externality is sufficiently large. In this case, liability should be preempted because experimentation needs to be encouraged. When the externality is small, instead, liability still has a role to play in addition to authorization regulation. In this case, not preempting liability allows the regulatory tools (approval and withdrawal) to be fixed at the first-best levels, liability being thus used to adjust the experimentation margin. Importantly, liability is not used to affect approval or withdrawal, but just to control experimentation.

Schwartzstein and Shleifer (2013) also highlight that positive externalities may make preemption desirable, in a model featuring two types of firms (safe and unsafe) that make an upfront investment in precaution. They show that preemption should hold only if unsafe firms generate positive social returns if they fail to take precautions. In this case, firms should be shielded from future liability if they meet the initial standards set by the regulator. We provide a general condition under which the preemption doctrine should not hold, based on a mechanism different from Schwartzstein and Shleifer (2013). We show that liability has a role to play to discourage socially excessive experimentation when ex ante and ex post precautions are optimally set by the regulator.

Innovative Industries and Digital Markets. We now turn to recent calls for increasing authorization regulations in the context of digital markets and innovative industries characterized by strong network effects. The Stigler Report proposes the creation of a new digital agency, in charge of monitoring the activities of firms. The report argues that "less disruptive ex post remedies require ongoing monitoring, which antitrust enforcers are not well-positioned to do. Handing that job off to a regulator might better serve consumers." According to this view, the agency could be in charge of setting ex ante rules on business practices for digital markets.

To address this application, our framework would need to be extended to take into account the interaction between firms and consumers. While we leave a full analysis to future research, based on our reduced-form model we expect the payoff for the firm in the two states to increase when competition is limited, thus resulting in a low positive externality in the good state. To fix ideas, in the limit case with $v_f = v_p^G$, by setting $L = v_p^G - v_p^B$ the firm's payoff in state $\theta = B$ is aligned with the social payoff; the first best is then achieved with full liability. In this sense, ex post competition policy is the optimal regulation if there is no positive externality.

Innovative industries, however, entail a positive externality in the good state. When $v_f < v_p^G$, our analysis predicts that some form of ex ante regulation should be introduced. By Proposition 5.(b), a lenient competition policy implementing a positive level of—but less than full—liability is called for when the positive externality is relatively small. In addition to highlighting the complementarity of ex post liability (competition policy) with authorization regulation, our analysis prescribes authorization regulation for innovative industries characterized by limited competition. However, for less mature industries, characterized by strong innovation externalities and more vigorous competition, the model prescribes that liability-based ex post competition policy should be lifted in favor of a lenient authorization regime.

Against the backdrop of the recent push toward ex ante regulation in digital markets, our model thus highlights the role of positive externalities. In the presence of large positive externalities, a move to ex ante regulation through the creation of a digital agency is particularly warranted for the purpose of preserving incentives to innovate. Current developments in the UK move in this direction. Following the Furman Review (2019), the Government has announced the creation of the Digital Markets Unit within the existing Competition and Markets Authority; the new unit has been designed as an agency with ex ante regulatory functions and with a deliberate focus on promoting further innovation (Competition and Markets Authority 2021).

The discussion in the Stigler Report highlights an additional dimension that could also be analyzed by extending our model. Ex ante rules also have the advantage of enhancing predictability in an environment characterized by large uncertainty. In our model, the firm knows the level of liability it will face if the state turns out to be bad. Adding uncertainty in the liability level might tilt the balance in favor of ex ante regulation. Recent events back the validity of this suggestion; the push toward ex ante regulation in digital markets has been driven from existing antitrust enforcement being criticized as uncertain and slow to react (Madiega 2020).

5 Conclusion

The paper proposes a tractable multi-stage model of experimentation and learning. Thanks to our simple formulation of incentives in terms of best replies, we are able not only to characterize the social planner solution, but also to analyze the strategic interaction between a firm with misaligned objectives and the regulator. Our positive and normative results shed light on the regulation of activities with uncertain externalities, with a special emphasis on safety.

As ex ante experimentation gives way to ex post learning in the field, the cost of information acquisition becomes more dependent on the state of the world. This feature is not specific to our two-stage formulation, but applies more generally to situations where experimentation progresses to more mature phases from the lab to the bench. Extending the techniques developed in this paper, it is possible to capture multiple phases of experimentation and the incentives induced.

This structure could be used to model the design of multiple testing phases, which play a key role in the regulation of clinical trials.

In the model, authorization regulation and liability are adjusted to preserve the experimentation incentives of the firm. In practice, other instruments such as subsidies can be used to encourage experimentation. If the regulator could subsidize the experimentation cost in the first phase without constraints, the first best could be achieved. In practice, however, there are some barriers. First of all, these subsidies need to be financed. Second, firms need to be incentivized to come forth with information about the value and cost of their inventions. Third, evaluators working for research funding organizations might have interests not perfectly aligned with social goals.

Throughout the analysis we posited that the regulator maximizes social welfare. In practice effective authorization regulation requires industry expertise, something difficult to develop in dynamic and innovative industries—the devil is in the details. As a complicating factor, specialized regulators can be easily captured by industry even in advanced democracies (Stigler 1971). More decentralized solutions, such as liability through private enforcement where victims can sue the injurers, may economize on information. We leave these important avenues to future work.

.

References

Adda J, Decker C, Ottaviani M (2020) P-hacking in Clinical Trials and How Incentives Shape the Distribution of Results Across Phases. *Proc. Natl. Acad. Sci. U.S.A.* 117(24):13386–13392.

Ahuja V, Alvarez C, Birge JR, Syverson C (2021) Enhancing Regulatory Decision-Making for Postmarket Drug Safety. *Management Sci.* forthcoming.

Angus DC (2020) Optimizing the Trade-off Between Learning and Doing in a Pandemic. *J. Amer. Med. Assoc.* 323(19):1895–1896.

Berniker S (2001) Spontaneous Reporting Systems: Achieving Less Spontaneity and More Reporting. Harvard Law School Third Year Paper.

Berry DA, Fristedt B (1985) *Bandit Problems* (New York: Chapman and Hall).

Bizzotto J, Rudiger J, Vigier A (2020) Testing, Disclosure and Approval. *J. Econom. Theory* 187:105002.

Bizzotto J, Rudiger J, Vigier, A (2021) Dynamic Persuasion with Outside Information *Amer. Econom. J. Microecon.* 13(1):179–194.

Bolton P, Harris C (1999) Strategic Experimentation. *Econometrica* 67(2):349–374.

Branco F, Sun M, Villas-Boas JM (2012). Optimal Search for Product Information. *Management Sci.* 58(11):2037–2056.

Cabral L (2021) Merger Policy in Digital Industries. *Inf. Econom. Policy* 54:100866.

Carpenter D (2004) Protection without Capture: Product Approval by a Politically Responsive, Learning Regulator. *Amer. Polit. Sci. Rev.* 98(4):613–631.

Chan J, Lizzeri A, Suen W, Yariv, L (2018) Deliberating Collective Decisions. *Rev. Econom. Stud.* 85(2):929–963.

Che YK, Mierendorff, K (2019) Optimal Dynamic Allocation of Attention. *Amer. Econom. Rev.* 109(8):2993–3029.

Competition and Markets Authority (2021) *The CMA's Digital Markets Strategy*. February 2021 Refresh.

Daughety AF, Reinganum, JF (2013), Economic Analysis of Products Liability: Theory. Chapter 3 in Arlen, J. (Ed.), *Research Handbook on the Economics of Torts*, Edward Elgar.

Dvoretsky A, Kiefer J, Wolfowitz J (1953) Sequential Decision Problems for Processes with Continuous Time Parameter: Testing Hypotheses. *Ann. Math. Stat.* 24(2):254–264.

Friehe T, Schulte E (2017) Uncertain Product Risk, Information Acquisition, and Product Liability. *Econom. Letters* 159:92–95.

Fudenberg D, Strack P, Strzalecki, T (2018) Speed, Accuracy, and the Optimal Timing of Choices. *Amer. Econom. Rev.* 108(12):3651–3684.

Furman J (2019) The Role of Data and Privacy in Competition. *Prepared for U.S. Congress Hearing Testimony on Online Platforms and Market Power.*

Galasso A, Luo H (2017) Tort Reform and Innovation. *J. Law Econom.* 60:385–412.

Galasso A, Luo H (2018). When Does Product Liability Risk Chill Innovation? Evidence From Medical Implants. *NBER Working Paper 25068*.

Garber S (2013) Economic Effects of Product Liability and Other Litigation Involving the Safety and Effectiveness of Pharmaceuticals. *RAND Corporation Institute for Civil Justice*.

Glaeser E, Shleifer, A (2003) The Rise of the Regulatory State. *J. Econom. Lit.* 41(2):401-425.

Green B, Taylor CR (2016) Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects. *Amer. Econom. Rev.* 106(12):3660–3699.

Grenadier SR, Malenko A, Malenko N (2016) Timing Decisions in Organizations: Communication and Authority in a Dynamic Environment. *Amer. Econom. Rev.* 106(9):2552–2581.

Guo Y (2016) Dynamic Delegation of Experimentation. *Amer. Econom. Rev.* 106(8):1969–2008.

Halac M, Kartik N, Liu Q (2016) Optimal Contracts for Experimentation. *Rev. Econom. Stud.* 83(3):1040–1091.

Hall BH, Graham S, Harhoff D, Mowery DC (2004) Prospects for Improving U.S. Patent Quality via Postgrant Opposition. *Innov. Policy Econom.* 4:115–143.

Hamada K (1976) Liability Rules and Income Distribution in Product Liability. *Amer. Econom. Rev.* 66(1):228–234.

Han L, Ball R, Pamer C, Altman R, Proestel S (2017) Development of an Automated Assessment Tool for MedWatch Reports in the FDA Adverse Event Reporting System. *J. Am. Med. Inform. Assoc.* 24(5):913–920.

Henderson JA, Twerski AD (1998) Achieving Consensus on Defective Product Design. *Cornell Law Rev.* 83(4):867–920.

Henry E, Ottaviani M (2019) Research and the Approval Process: The Organization of Persuasion. *Amer. Econom. Rev.* 109(3):911-955.

Hua X, Spier KE (2020) Product Safety, Contracts, and Liability. *RAND J. Econom.* 51(1):233–259.

Inderst R, Ottaviani M (2013) Sales Talk, Cancellation Terms and the Role of Consumer Protection. *Rev. Econom. Stud.* 80(3):1002–1026.

Jaffe B, Lerner J (2007) *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do About It* (Princeton University Press).

Jakubiak J (1997) Maintaining Air Safety at Less Cost: A Plan for Replacing FAA Safety Regulations with Strict Liability. *Cornell J. Law Public Policy* 6(2):421–440.

Kamenica E, Gentzkow, M (2011) Bayesian Persuasion. *Amer. Econom. Rev.* 101(6):2590–615.

Kartik N, Ottaviani M, Squintani F (2007) Credulity, Lies, and Costly Talk. *J. Econom. Theory* 134(1):93–116.

Kolstad CD, Ulen TS, Johnson GV (1990) Ex Post Liability for Harm vs. Ex Ante Safety Regulation: Substitutes or Complements? *Amer. Econom. Rev.* 80(4):888–901.

Konert A (2019) Aviation Accidents Involving Boeing 737 MAX: Legal Consequences. *Ius Novum* 3:119–133.

Lybecker KM, Watkins L (2015) Liability Risk in the Pharmaceutical Industry: Tort Law in the US and UK. *Soc. Sci. J.* 52(4):433–448.

Madiega T (2020) Regulating Digital Gatekeepers: Background on the Future Digital Markets Act. European Parliamentary Research Service, PE 659.397.

Marino AM (1997). A Model of Product Recalls with Asymmetric Information. *J. Reg. Econom.* 12(3):245–265.

McClellan A (2017) Experimentation and Approval Mechanisms. New York University Working Paper.

Morris S, Strack P (2017) The Wald Problem and the Equivalence of Sequential Sampling and Static Information Costs. SSRN Working Paper.

Moscarini G, Smith L (2001). The Optimal Level of Experimentation. *Econometrica* 69(6):1629–1644.

Office of Inspector General (2015). *Audit Report: FAA Lacks an Effective Staffing Model and Risk-Based Oversight Process for Organization Designation Authorization*. Federal Aviation Administration, Report Number AV-2016-001.

Oi WY (1973) The Economics of Product Safety. *Bell J. Econom. Management Sci.* 4(1):3–28.

Orlov D, Skrzypacz A, Zryumov P (2020) Persuading the Principal to Wait. *J. Political Econom.* 128(7):2542–2578.

Ottaviani M, Wickelgren AL (2009) Approval Regulation and Learning, with Application to Timing of Merger Control. Northwestern University Working Paper.

Ottaviani M, Wickelgren AL (2011) Ex Ante or Ex Post Competition Policy? A Progress Report. *Int. J. Ind. Organ.* 29(2):356–359.

Polinsky M, Shavell S (2010) The Uneasy Case for Product Liability. *Harv. L. Rev.* 123(6):1437–1493.

Pomatto L, Strack P, and Tamuz O (2019) The Cost of Information. SSRN Working Paper.

Rey P (2003) Towards a Theory of Competition Policy. In: Dewatripont M, Hansen LP, Turnovsky SJ (Eds.), *Advances in Economics and Econometrics: Theory and Applications* (Eighth World Congress Cambridge University Press, Cambridge).

Rupp NG, Taylor CR (2002) Who Initiates Recalls and Who Cares? Evidence from the Automobile Industry. *J. Ind. Econom.* 50(2):123-149.

Schmitz PW (2000) On the Joint Use of Liability and Safety Regulation. *Int. Rev. Law Econom.* 20(3):371–382.

Schwartzstein J, Shleifer A (2013) An Activity-Generating Theory of Regulation. *J. Law Econom.* 56(1):1–38.

Scott-Morton F, Bouvier P, Ezrachi A, Jullien B, Katz, R, Kimmelman G, Melamed D, Morgenstern, J (2019) Committee for the Study of Digital Platforms. *Report, Stigler Center for the Study of the Economy and the State.*

Sharkey CM (2008) Products Liability Preemption: An Institutional Approach. *Geo. Wash. Law Rev.* 76(3):449–521.

Shavell S (1984) Liability for Harm Versus Regulation of Safety. *J. Legal Stud.* 13(2):357–374.

Spier KE (2011) Product Safety, Buybacks, and the Post-Sale Duty to Warn. *J. Law Econom. Organ.*, 27(3): 515–539.

Stigler GJ (1971) The Theory of Economic Regulation. *Bell J. Econom. Management Sci.* 2(1):3–21.

Stokey N (2009) *The Economics of Inaction* (Princeton University Press, Princeton).

Strulovici B (2010) Learning While Voting: Determinants of Collective Experimentation. *Econometrica* 78(3):933–971.

Van Norman GA (2016). Drugs, Devices, and the FDA: Part 1: An Overview of Approval Processes for Drugs. *J. Amer. Coll. Card. Basic Trans. Sci.* 1(3):170–179.

Viscusi W, Moore M (1993) Product Liability, Research and Development, and Innovation. *J. Political Econom.* 101(1):161–84.

Viscusi WK, Vernon JM, Harrington JE (1995) *Economics of Regulation and Antitrust* (MIT Press).

Wald A (1945) Sequential Tests of Statistical Hypotheses. *Ann. Math. Stat.* 16(2):117–186.

Welling L (1991) A Theory of Voluntary Recalls and Product Liability. *South. Econom. J.* 57(4):1092–1111.

Zhong W (2019) Optimal Dynamic Information Acquisition. Columbia University Working Paper.

Zuckerman DM, Brown P, Nissen SE (2011) Medical Device Recalls and the FDA Approval Process. *Arch. Inter. Med.*, 171(11):1006–1011.

A Appendix A: Construction of Best Replies

This appendix presents some foundational results for the construction and comparative statics of the best replies. These results are used in Appendix B to prove the propositions in the main text. We report the results for an arbitrary player j who obtains payoffs $v_j^G > 0$ in the good state and $v_j^B < 0$ in the bad state. In the first-best benchmark player j corresponds to the planner. In the strategic analysis player $j = f$ corresponds to the firm facing a liability rate $L > v_f$. Appendix C proves the results reported in this appendix.

A0. Log-odds Parametrization. To facilitate exposition, we derive results using the following log-odds re-parametrization of beliefs

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

In this log-odds space, experimentation, approval and withdrawal standards are defined accordingly as $\mathsf{s} = \log \frac{s}{1-s}$, $\mathsf{S} = \log \frac{S}{1-S}$, and $\mathsf{z} = \log \frac{z}{1-z}$. We state the results in the regular belief space, but use the log-odds parametrization for the proofs.

A1. Hitting Probabilities. We start by reporting some preliminary results about hitting probabilities. The expected unconditional hitting probabilities in the first phase are defined by

$$\Psi^{\mathrm{I}}(\sigma) = \frac{e^{\sigma}}{1+e^{\sigma}} \Psi^{\mathrm{I}}(\sigma, G) + \frac{1}{1+e^{\sigma}} \Psi^{\mathrm{I}}(\sigma, B), \quad \psi^{\mathrm{I}}(\sigma) = \frac{e^{\sigma}}{1+e^{\sigma}} \psi^{\mathrm{I}}(\sigma, G) + \frac{1}{1+e^{\sigma}} \psi^{\mathrm{I}}(\sigma, B).$$

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

$$\begin{aligned} \Psi^{\mathrm{I}}(\sigma, G) &= \tfrac{e^{-R_1(\sigma-\mathsf{s})} - e^{-R_2(\sigma-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})} - e^{-R_2(\mathsf{S}-\mathsf{s})}}, \quad \Psi^{\mathrm{I}}(\sigma, B) = e^{\sigma-\mathsf{S}} \Psi^{\mathrm{I}}(\sigma, G), \\ \psi^{\mathrm{I}}(\sigma, G) &= \tfrac{e^{R_2(\mathsf{S}-\sigma)} - e^{R_1(\mathsf{S}-\sigma)}}{e^{R_2(\mathsf{S}-\mathsf{s})} - e^{R_1(\mathsf{S}-\mathsf{s})}}, \quad \psi^{\mathrm{I}}(\sigma, B) = e^{\sigma-\mathsf{s}} \psi^{\mathrm{I}}(\sigma, G), \end{aligned}$$

where $R_1 = \frac{1}{2}\left(1 - \sqrt{1 + \frac{4r}{\mu_{\mathrm{I}}}}\right)$ and $R_2 = \frac{1}{2}\left(1 + \sqrt{1 + \frac{4r}{\mu_{\mathrm{I}}}}\right)$ satisfy $R_1 < 0$, $R_2 > 0$, and $R_1 + R_2 = 1$.

For the second phase, given withdrawal standard z, the expected unconditional hitting probability is

$$\psi^{\mathrm{II}}(\sigma) = \frac{e^{\sigma}}{1+e^{\sigma}} \psi^{\mathrm{II}}(\sigma, G, \mathsf{z}) + \frac{1}{1+e^{\sigma}} \psi^{\mathrm{II}}(\sigma, B, \mathsf{z})$$

where

$$\psi^{\mathrm{II}}(\sigma, G, \mathsf{z}) = e^{-r_2(\sigma-\mathsf{z})} \quad \text{and} \quad \psi^{\mathrm{II}}(\sigma, B, \mathsf{z}) = e^{r_1(\sigma-\mathsf{z})},$$

with $r_1 = \frac{1}{2}\left(1 - \sqrt{1 + \frac{4r}{\mu_{\mathrm{II}}}}\right)$ and $r_2 = \frac{1}{2}\left(1 + \sqrt{1 + \frac{4r}{\mu_{\mathrm{II}}}}\right)$ such that $r_1 < 0$, $r_2 > 0$, $r_1 + r_2 = 1$.

Lemma 1 *In the first phase, the expected conditional hitting probabilities Ψ^{I} and ψ^{I} satisfy*

$$\begin{aligned} (1) \quad \frac{\partial \Psi^{\mathrm{I}}(\sigma, G)}{\partial \mathsf{S}} &= f \Psi^{\mathrm{I}}(\sigma, G) < 0, \quad (2) \quad \frac{\partial \Psi^{\mathrm{I}}(\sigma, G)}{\partial \mathsf{s}} = a \psi^{\mathrm{I}}(\sigma, G) < 0, \\ (3) \quad \frac{\partial \psi^{\mathrm{I}}(\sigma, G)}{\partial \mathsf{S}} &= g \Psi^{\mathrm{I}}(\sigma, G) > 0, \quad (4) \quad \frac{\partial \psi^{\mathrm{I}}(\sigma, G)}{\partial \mathsf{s}} = b \psi^{\mathrm{I}}(\sigma, G) > 0, \end{aligned}$$

where

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

In the second phase, the expected conditional hitting probabilities $\psi^{\mathrm{II}}(\sigma, G, \mathsf{z})$ *and* $\psi^{\mathrm{II}}(\sigma, B, \mathsf{z})$ *satisfy*

$$
\begin{aligned}
(5) \quad \frac{\partial \psi^{\mathrm{II}}(\sigma, G, \mathsf{z})}{\partial \sigma} &= -\frac{\partial \psi^{\mathrm{II}}(\sigma, G, \mathsf{z})}{\partial \mathsf{z}} = -r_2 \psi^{\mathrm{II}}(\sigma, G, \mathsf{z}) > 0, \\
(6) \quad \frac{\partial \psi^{\mathrm{II}}(\sigma, B, \mathsf{z})}{\partial \sigma} &= -\frac{\partial \psi^{\mathrm{II}}(\sigma, B, \mathsf{z})}{\partial \mathsf{z}} = r_1 \psi^{\mathrm{II}}(\sigma, B, \mathsf{z}) > 0.
\end{aligned}
$$

A2. Expected Payoff Functions and Best Replies. In log-odds, from equation (1) in the main text, the expected payoff of player j in the ex post phase is given by:

$$
u_j^{\mathrm{II}}(\mathsf{S}) = \frac{e^{\mathsf{S}}}{1+e^{\mathsf{S}}} \left(\frac{v_j^G}{r} \right) (1 - \psi^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})) + \frac{1}{1+e^{\mathsf{S}}} \left(\frac{v_j^B}{r} \right) (1 - \psi^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})), \tag{8}
$$

where

$$
\psi^{\mathrm{II}}(\mathsf{S}, \mathsf{z}) = \frac{e^{\mathsf{S}}}{1+e^{\mathsf{S}}} \psi^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}) + \frac{1}{1+e^{\mathsf{S}}} \psi^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}).
$$

Similarly the expected payoff of player j in the ex ante phase

$$
u_j^{\mathrm{I}}(\sigma) = \psi^{\mathrm{I}}(\sigma) u_j^{\mathrm{II}}(\mathsf{s}) + \Psi^{\mathrm{I}}(\sigma) u_j^{\mathrm{II}}(\mathsf{S}) - \left[1 - \psi^{\mathrm{I}}(\sigma) - \Psi^{\mathrm{I}}(\sigma)\right] \frac{c}{r}, \tag{9}
$$

where the term $u_j^{\mathrm{II}}(s) \equiv 0$ whenever $s < z$. Using Lemma 1 we can rewrite (9) as

$$
\begin{aligned}
u_j^{\mathrm{I}}(\sigma) &= -\frac{c}{r} + \frac{e^{\sigma}}{1+e^{\sigma}} \Psi^{\mathrm{I}}(\sigma, G) \left[u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}) + e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}) + (1 + e^{-\mathsf{S}}) \frac{c}{r} \right] \\
&\quad + \frac{1}{1+e^{\sigma}} \psi^{\mathrm{I}}(\sigma, G) \left[u_j^{\mathrm{II}}(\mathsf{s}, G, \mathsf{z}) + e^{-\mathsf{s}} u_j^{\mathrm{II}}(\mathsf{s}, B, \mathsf{z}) + (1 + e^{-\mathsf{s}}) \frac{c}{r} \right]
\end{aligned} \tag{10}
$$

where the term $u_j^{\mathrm{II}}(\mathsf{s}, G, \mathsf{z}) + e^{-\mathsf{s}} u_j^{\mathrm{II}}(\mathsf{s}, B, \mathsf{z})$ is zero whenever $\mathsf{s} < \mathsf{z}$ and $u_j^{\mathrm{II}}(\mathsf{S}, \theta, \mathsf{z}) = \frac{v_j^{\theta}}{r}(1 - \psi^{\mathrm{II}}(\mathsf{S}, \theta, \mathsf{z}))$ for $\theta \in \{G, B\}$.

In the ex post phase, the withdrawal standard is optimally chosen to maximize (8). Using the hitting probabilities from Section A1, the solution can be expressed in closed form as

$$
\mathsf{z}^* = -\ln \frac{v_j^G r_2}{v_j^B r_1}. \tag{11}
$$

In the ex ante phase, as explained in the main text, the optimal rejection and approval standards $(\mathsf{s}^*, \mathsf{S}^*)$ are at the intersection of the upper and lower best replies, with withdrawal in the ex post phase set at z^*. The upper and lower best replies are characterized by equations (5) and (4) in the main text.

The following result characterizes the shape of the best replies in the regular belief space, as illustrated in Figures 6 and 7:

Lemma 2 *Given initial belief q, if player j withdraws at z^*, in the ex ante phase:*

(a) *the lower best reply $s = b_j(S)$ (i) initially grows along the diagonal $b_j(S) = S$ for $S \in [0, z^*]$, (ii) decreases in S for $S \in [z^*, S^*]$ and (iii) increases in S for $S \in [S^*, 1]$ to reach 1 at the limit as $S \to 1$.*

(b) *the upper best reply $S = B_j(s)$ (i) is constant at $B_j(s) = S^*$ for $s \in [0, s^*]$, (ii) decreasing in s for $s \in [s^*, \bar{s}]$ and (iii) grows along the diagonal $B_j(s) = s$ for $s \in (\bar{s}, 1]$.*

In log-odds the lower best reply $b_j(\mathsf{S})$, when interior, is implicitly defined from the first order condition of (10) with respect to s which using results in Lemma 1 becomes[45]

$$\begin{aligned}\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{s}} &= \frac{e^\sigma}{1+e^\sigma}\psi^{\mathrm{I}}(\sigma, G)\Big[a\Big[u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}}u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) + (1+e^{-\mathsf{S}})\frac{c}{r}\Big] \\ &\quad + b(1+e^{-\mathsf{s}})\frac{c}{r} - e^{-\mathsf{s}}\frac{c}{r}\Big] = 0. \end{aligned} \tag{12}$$

Similarly, in log-odds the upper best reply $B_j(\mathsf{s})$, when interior, is implicitly defined by the first order condition of (10) with respect to S. By Lemma 1, the upper best reply thus satisfies

$$\begin{aligned}\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{S}} = \frac{e^\sigma}{1+e^\sigma}\Psi^{\mathrm{I}}(\sigma, G)\Big[f\Big[u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}}u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) + (1+e^{-\mathsf{S}})\frac{c}{r}\Big] \\ - e^{-\mathsf{S}}\Big[u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) + \frac{c}{r}\Big] + g(1+e^{-\mathsf{s}})\frac{c}{r} + \Big[\frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*)}{\partial \mathsf{S}} + e^{-\mathsf{S}}\frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*)}{\partial \mathsf{S}}\Big]\Big] = 0,\end{aligned} \tag{13}$$

where the last term in parentheses is the postponement effect, which is zero at $\mathsf{z} = \mathsf{z}^*$.[46]

A3. Shifts of Best Replies. Having described the shape of the best replies, we now characterize how they shift with changes in v_j^G and v_j^B, as also illustrated in Figures 6 and 7:

Lemma 3 *As v_j^G increases:*

(a) *the withdrawal standard z^* decreases,*

(b) *the lower best reply $s = b_j(S)$ shifts to the left for $S \in [z^*, 1]$,*

[45]Note that the term $\left(u_j^{\mathrm{II}}(\mathsf{s}, G, \mathsf{z}) + e^{-\mathsf{s}}u_j^{\mathrm{II}}(\mathsf{s}, B, \mathsf{z})\right)$ in (10) shall not be considered when computing $\frac{\partial u^{\mathrm{I}}}{\partial \mathsf{s}}$ since a player choosing s while taking S as given has only the option to reject and obtain zero.

[46]Note that in equation (13) we are considering $\mathsf{s} < \mathsf{z}^*$. If $\mathsf{s} > \mathsf{z}^*$ we would have the additional term $g\left(u_j^{\mathrm{II}}(\mathsf{s}, G, \mathsf{z}^*) + e^{-\mathsf{s}}u_j^{\mathrm{II}}(\mathsf{s}, B, \mathsf{z}^*)\right)$.

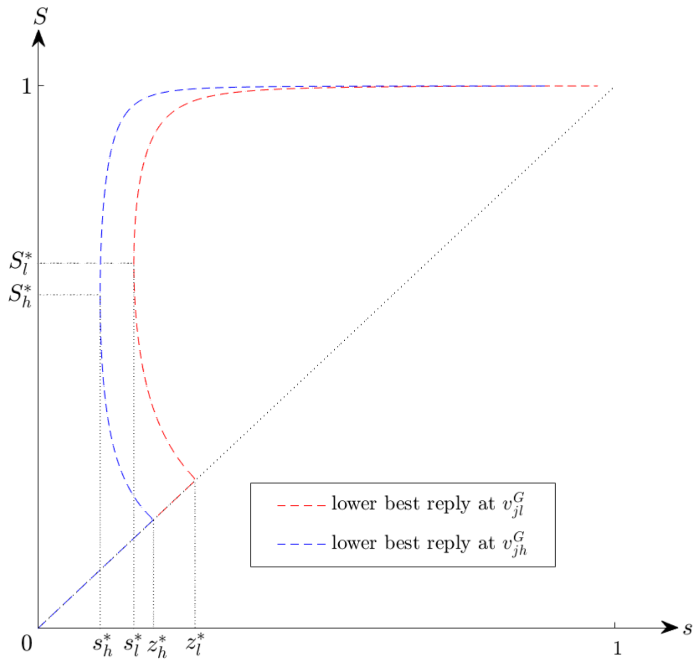

Figure 6: Lower best reply as v_j^G increases.

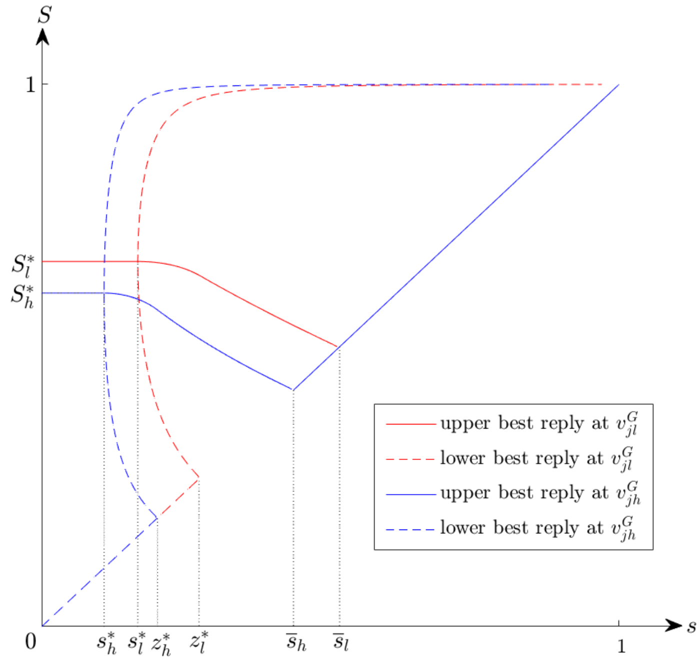

Figure 7: Upper best reply as v_j^G increases.

(c) *the upper best reply $S = B_j(s)$ shifts down for $s \in [s^*, \bar{s}]$.*

Lemma 4 *As v_j^B increases:*

(a) *the withdrawal standard z^* decreases,*

(b) *the lower best reply $s = b_j(S)$ shifts to the left for $S \in [z^*, 1]$,*

(c) *the upper best reply $S = B_j(s)$ shifts down for $s \in [s^*, \bar{s}]$.*

Parameterizing payoffs by α with $v_j^B = \alpha \bar{v}_j^B$ and $v_j^G = \alpha \bar{v}_j^G$, we characterize the impact of an increase in the payoffs in the bad and good state while keeping their ratio fixed:

Lemma 5 *As α increases:*

(a) *the lower best reply $s = b_j(S)$ shifts to the left for $S \in [z^*, 1]$,*

(b) *the upper best reply $S = B_j(s)$ shifts upwards for $s \in [s^*, \bar{s}]$.*

The comparative statics of the best replies with respect to z is characterized as follows:

Lemma 6 *Let z be any given withdrawal standard and S be any given approval standard. The lower best reply $s = b_j(S)$ and the upper best reply $S = B_j(s)$, when interior, are increasing in z if and only if $z > z^*$.*

B Appendix B: Proofs of Results in Main Text

Proof of Proposition 1

Step 1. *The planner optimally chooses standards (s^*, S^*, z^*) independent from the initial belief q.*

In the ex post phase, given any belief σ, the optimal withdrawal z^* (which solves $\frac{\partial u_p^{\mathrm{II}}(\mathsf{s})}{\partial \mathsf{z}} = 0$) is given by $\mathsf{z}^* = -\ln \frac{r_2 v_p^G}{r_1 v_p^B}$ and does not depend on σ, as explained in Appendix A. Given z^* and any initial belief σ, the two optimal thresholds $(\mathsf{s}^*, \mathsf{S}^*)$ in the ex ante phase are determined by the intersection of the lower and upper best replies $b_p(\mathsf{S})$ and $B_p(\mathsf{s})$.

The lower best reply $b_p(\mathsf{S})$ for $\mathsf{S} > \bar{\mathsf{z}}$, is implicitly defined by the first-order condition $\frac{\partial u_p^{\mathrm{I}}(\sigma)}{\partial \mathsf{s}} = 0$. Taking expression (12) for player $j = p$, $b_p(\mathsf{S})$ is characterized by

$$a\left[u_p^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}} u_p^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) + (1 + e^{-\mathsf{S}})\frac{c}{r}\right] + b(1 + e^{-\mathsf{s}})\frac{c}{r} - e^{-\mathsf{s}}\frac{c}{r} = 0, \tag{14}$$

which implies that $b_p(\mathsf{S})$ does not depend on σ. Similarly, the upper best reply $B_p(\mathsf{s})$, when interior, is implicitly defined by the first-order condition $\frac{\partial u_p^{\mathrm{I}}(\sigma)}{\partial \mathsf{S}} = 0$. Taking expression (13) at $\mathsf{z} = \mathsf{z}^*$ for player p, $B_p(\mathsf{s})$ is characterized by

$$f\left[u_p^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}} u_p^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) + (1 + e^{-\mathsf{S}})\frac{c}{r}\right] - e^{-S}\left[u_p^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) + \frac{c}{r}\right] + g(1 + e^{-\mathsf{s}})\frac{c}{r} = 0,$$

where the last term in expression (13), which corresponds to the postponement effect, is zero for the optimally chosen withdrawal $\mathsf{z} = \mathsf{z}^*$. This implies that $B_p(\mathsf{s})$ does not depend on σ.

Step 2. *The optimal standards are such that* $\mathsf{s}^* \le \mathsf{z}^* \le \mathsf{S}^*$.

Suppose that $\mathsf{z}^* > \mathsf{S}^*$. Approval would lead to immediate withdrawal (if $\mathsf{S}^* \le z^*$), with no experimentation by the firm, thus leading to a contradiction. Similarly, suppose that $\mathsf{z}^* < \mathsf{s}^*$, then, for beliefs $\sigma \in (\mathsf{z}^*, \mathsf{s}^*)$, the planner would be better off approving rather than rejecting, since that would yield positive payoffs. This leads to a contradiction.

Proof of Proposition 2

Step 1. *If $L = \overline{L} \equiv -e^B$, firm f withdraws too early $z^* < z_{ff}$, rejects too early $s^* < s_{ff}$, and approves too late $S^* < S_{ff}$.*

In the case of liability, the firm f solves a stand alone problem with payoffs decreased by L in the bad state. For $L = \overline{L} = -e^B$ the firm has socially insufficient incentives in the good state $v_f < v_p^G$ and perfectly aligned incentives in the bad state $v_f - \overline{L} = v_p^B$. Furthermore, from Lemma 3 we know that for any player j as v_j^G increases the withdrawal standard z^* decreases, the upper best reply $B_j(s)$ shifts down, and the lower best reply $b_j(S)$ shifts to the left. This implies that the firm, which faces the same problem as the planner but with lower payoff v_j^G in the good state, withdraws too early, $z^* < z_{ff}$. This also implies that $B_f(s)$, when interior, lies everywhere above $B_p(s)$ so that the firm approves too late $S^* < S_{ff}$ and $b_f(S)$, when interior, is everywhere to the right of $b_p(S)$ so that the firm abandons too early, $s^* < s_{ff}$.

Step 2. *The optimal liability L^*_{ff} is interior with $v_f < L^*_{ff} < \overline{L} \equiv -e^B$.*

If $L \leq v_f$, firm f has positive payoffs in both states and hence no incentives to conduct costly experimentation nor to withdraw: $s_{ff} = z_{ff} = S_{ff} = 0$. The planner thus chooses $L^*_{ff} > v_f$. From step 1 we know that for $L = \overline{L}$, the firm withdraws too early $z^* < z_{ff}$, rejects too early $s^* < s_{ff}$, and approves too late $S^* < S_{ff}$. Reducing L increases the firm's payoff in the bad state $v_f - L$. By Lemma 4, as v^B_j increases, z^* decreases, the upper best reply $B_j(s)$ shifts down, and the lower best reply $b_j(S)$ shifts to the left. Thus, when L is reduced, all three standards s_{ff}, z_{ff} and S_{ff} decrease. Marginally decreasing L from $L = \overline{L}$ moves all three standard toward the socially optimal levels, so that it optimal for the planner to set $L^*_{ff} < \overline{L}$, *proving result (a).*

Step 3. *There exists $\hat{L} < \overline{L}$ such that $z_{ff}(\hat{L}) = z^*$, $s^* < s_{ff}(\hat{L})$ and $S_{ff}(\hat{L}) < S^*$.*

To induce the firm to withdraw at $\mathsf{z}_{ff}(L) = \mathsf{z}^*$, the liability L must be set such that the ratio of payoffs in the good and bad state is the same for the firm f as for the planner p: $\frac{v_f - \hat{L}}{v_f} = \frac{v^B_p}{v^G_p}$. This gives $\hat{L} = \frac{v_f(v^G_p - v^B_p)}{v^G_p}$. Given that the firm does not obtain the full social benefits in the good state ($v_f < v^G_p$), there exists a factor $\alpha < 1$ such that $v^B_f = \alpha v^B_p$ and $v^G_f = \alpha v^G_p$. By Lemma 5, the lower best reply of the planner is shifted to the left compared to the lower best reply of the firm for $L = \hat{L}$ and the upper best reply of the planner is shifted upward compared to the upper best reply of the firm for $L = \hat{L}$, as also shown in Figure 2. Thus, $s^* < s_{ff}(\hat{L}) < S_{ff}(\hat{L}) < S^*$.

Step 4. *The optimal liability level L^*_{ff} is such that $s^*_{ff} < s_{ff} < S_{ff} < S^*_{ff}$.*

Suppose $L^*_{ff} > \hat{L}$, which implies that firm f's payoff in the bad state is now lower than before: $v_f - L^*_{ff} < v_f - \hat{L}$. By Lemma 4 we thus have $s_{ff}(\hat{L}) < s_{ff}(L^*_{ff})$. Since $s^* < s_{ff}(\hat{L})$ by step 3, this implies that $s^* < s_{ff}(L^*_{ff})$. Lemma 4 also implies that $z_{ff}(\hat{L}) < z_{ff}(L^*_{ff})$ and given that $z_{ff}(\hat{L}) = z^*$, this implies that $z^* < z_{ff}(L^*_{ff})$. By way of contradiction suppose then that $S_{ff}(L^*_{ff}) > S^*$. Reducing the liability level from L^*_{ff} would then move all standards closer to their optimal values, contradicting the optimality of L^*_{ff}. Hence it must be that $S_{ff}(L^*_{ff}) < S^*$. Similarly, if $L^*_{ff} < \hat{L}$, assuming $s_{ff}(L^*_{ff}) < s^*$ would lead to a contradiction. *This proves result (b).*

Step 5. *The optimal liability L^*_{ff} is a function of the initial belief q. Moreover, $L^*_{ff} < \hat{L}$ if and only if $q \in [\underline{q}, \hat{q})$ and $L^*_{ff} > \hat{L}$ if and only if $q \in (\hat{q}, \overline{q}]$.*

If $q \leq s^*$ or $q \geq S^*$, it is optimal for the planner to commit to a liability rate L such that firm f immediately abandons in the first case or immediately approves in the second case.

Now consider $q \in (s^*, S^*)$, equivalent to $\sigma \in (\mathsf{s}^*, \mathsf{S}^*)$ in the log-odds space. The optimal liability level L^*_{ff} is obtained by taking the first order condition with respect to L of expression (10) for $j = p$

$$\frac{\partial u^{\mathrm{I}}_p(\sigma)}{\partial L} = \frac{\partial u^{\mathrm{I}}_p(\sigma)}{\partial \mathsf{z}} \frac{\partial \mathsf{z}_{ff}}{\partial L} + \frac{\partial u^{\mathrm{I}}_p(\sigma)}{\partial \mathsf{s}} \frac{\partial \mathsf{s}_{ff}}{\partial L} + \frac{\partial u^{\mathrm{I}}_p(\sigma)}{\partial \mathsf{S}} \frac{\partial \mathsf{S}_{ff}}{\partial L}. \tag{15}$$

For $L = \hat{L}$ the first term is zero since $z_{ff}(\hat{L}) = z^*$. The second term is negative since at $L = \hat{L}$ the planner's lower best reply is to the left of the firm f 's best reply implying $\frac{\partial u^{\mathrm{I}}_p(\sigma)}{\partial \mathsf{s}} < 0$ and

$\frac{\partial \mathsf{s}_{ff}}{\partial L} > 0$, given that by Lemma 4 the lower best reply shifts to the left as the payoff in the bad state increases. By a similar argument, the last term is positive. Overall at $L = \hat{L}$ the sign of (15) is determined by the initial belief sigma σ . From equations (12) and (13) with $j = p$ it is immediate to check that $\frac{\partial u_p^{\mathrm{I}}(\sigma)}{\partial \mathsf{s}}$ decreases in σ and $\frac{\partial u_p^{\mathrm{I}}(\sigma)}{\partial \mathsf{S}}$ increases in σ. As $\sigma \to \mathsf{s}^*$ the marginal value of increasing approval $\frac{\partial u_p^{\mathrm{I}}(\sigma)}{\partial \mathsf{S}} \to 0$ implying that (15) becomes negative. Similarly, as $\sigma \to \mathsf{S}^*$ the marginal value of increasing rejections $\frac{\partial u_p^{\mathrm{I}}(\sigma)}{\partial \mathsf{s}} \to 0$, implying that (15) becomes positive.

By continuity there exists $\hat{\sigma} \in (\mathsf{s}^*, \mathsf{S}^*)$ such that at $L = \hat{L}$ (15) is zero and thus $L^*_{ff} = \hat{L}$. For $\sigma < \hat{\sigma}$ the first order condition (15) is negative and thus $L^*_{ff} < \hat{L}$ so as to provide incentives for experimentation. For $\sigma > \hat{\sigma}$ the first order condition (15) is positive and thus $L^*_{ff} > \hat{L}$ so as to delay approval. *This proves result (c).*

Proof of Proposition 3

(a) For any value of c, under withdrawal regulation the firm ex post value is zero for any $q \in (0, z)$ and positive for $q > z$. In particular, at any $z > 0$ the smooth pasting condition of the firm ex post value with the zero horizontal line cannot hold. This is because from the firm perspective it is never optimal to withdraw. This generates a kink in the ex post value at z which ensures that experimentation is valuable for any c.

(b) The payoff in the bad state under the liability regime $L(\mathrm{z})$ is clearly lower than under withdrawal regulation at z, because $v_f - L(\mathrm{z}) < v_f$. Lemma 4 (comparative statics of the best replies for changes in v_j^B) is derived in the case where player j chooses withdrawal optimally and thus does not apply directly for the case of withdrawal regulation (when z is set by the planner). However, the proof of Lemma 4.(b) applies and the lower best reply $b_f(S)$ is shifted to the left under withdrawal regulation, so that $s_{fp}(z) < s_{ff}(L(z))$.

(c) As explained in the main text, the proof of Lemma 4.(c) no longer applies in the case of withdrawal regulation because of the postponement effect. By implicit differentiation and concavity, the sign of $\frac{\partial B_f(s)}{\partial v_f^B}$ is determined by $\frac{\partial u_f^{\mathrm{I}}(\sigma)}{\partial \mathsf{S} \partial v_j^B}$. From (13) for $j = f$

$$\frac{\partial u_f^{\mathrm{I}}(\sigma)}{\partial \mathsf{S} \partial v_f^B} = \frac{e^{\sigma}}{1+e^{\sigma}} \Psi^{\mathrm{I}}(\sigma, G) e^{-\mathsf{S}} \left[f \frac{\partial u_f^{\mathrm{II}}(\mathsf{S}, B, \mathrm{z})}{\partial v_f^B} - \frac{\partial u_f^{\mathrm{II}}(\mathsf{S}, B, \mathrm{z})}{\partial v_f^B} + \frac{\partial u_f^{\mathrm{II}}(\mathsf{S}, B, \mathrm{z})}{\partial \mathsf{S} \partial v_f^B} \right].$$

The first term in brackets is the opportunity cost of delaying approval, negative by $f < 0$. The second term is the value of information and is also negative. Both these terms appear in the proof of Lemma 4.(c).[47] The third term, which captures the postponement effect, is positive.

[47] The indirect effect, relevant in the proof of Lemma 4.(c), does not appear in our present context. Indeed this indirect effect corresponds to changes in z following shifts in v_f^B. In our current context, z remains fixed.

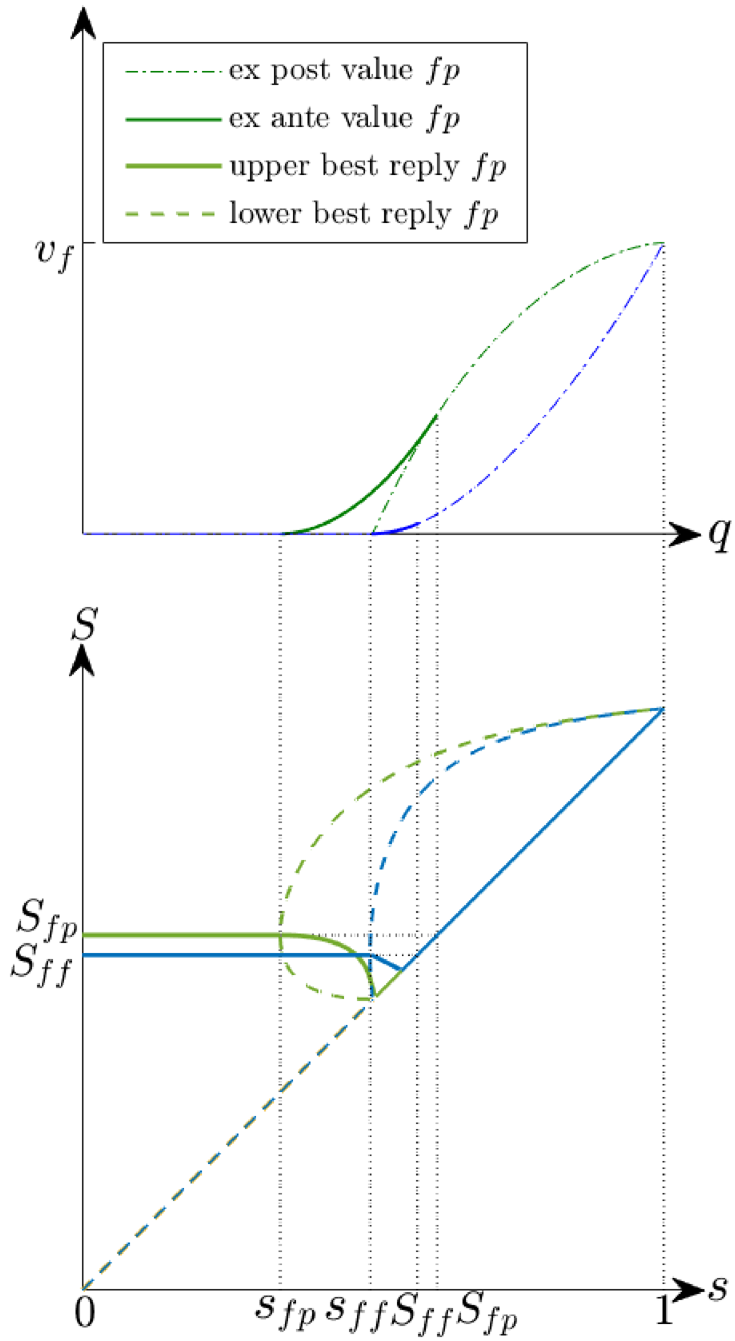

Figure 8: Value functions (top panel) and best replies (bottom panel) when the postponement effect dominates.

A sufficient condition that guarantees that $\frac{\partial u_f^{\mathrm{I}}(\sigma)}{\partial \mathsf{S} \partial v_f^B} < 0$ is thus

$$-\frac{\partial u_f^{\mathrm{II}}(\mathsf{S},B,\mathsf{z})}{\partial v_f^B} + \frac{\partial u_f^{\mathrm{II}}(\mathsf{S},B,\mathsf{z})}{\partial \mathsf{S} \partial v_f^B} < 0$$

where $u_f^{\mathrm{II}}(\mathsf{S},B,\mathsf{z}) = v_f^B \left[1 - \psi^{\mathrm{II}}(\mathsf{S},B,\mathsf{z})\right]/r$. By Lemma 1 this condition is equivalent to

$$-(1-\psi^{\mathrm{II}}(\mathsf{S},B,\mathsf{z})) - \frac{\partial \psi^{\mathrm{II}}(\mathsf{S},B,\mathsf{z})}{\partial \mathsf{S}} = -\left[1-\psi^{\mathrm{II}}(\mathsf{S},B,\mathsf{z})\right] - r_1 \psi^{\mathrm{II}}(\mathsf{S},B,\mathsf{z}) < 0$$

where $r_1 < 0$ is defined in Appendix A. Thus, defining $\bar{\psi} = \frac{1}{1-r_1}$, we conclude that $\psi^{\mathrm{II}}(\mathsf{S},B,\mathsf{z}) < \bar{\psi}$ implies $\frac{\partial B_f(s)}{\partial v_f^B} < 0$, so that $\mathsf{S}_{fp}(\mathsf{z}) < \mathsf{S}_{ff}(\tilde{L}(\mathsf{z}))$. This condition holds if for instance c is sufficiently low as it ensures that $S_{fp}(\mathsf{z})$ is set sufficiently far from z so that $\psi^{\mathrm{II}}(\mathsf{S}_{fp}(\mathsf{z}),B,\mathsf{z})$ is low.

Proof of Proposition 4

Under authorization regulation, the planner controls the approval and withdrawal standards and the firm chooses the experimentation standards. Thus the key ingredient of the proof is the comparative statics of the lower best reply.

Suppose both regulatory tools are chosen at their first-best levels, $S_{pp} = S^*$ and $z_{pp} = z^*$:

- If $e^G = 0$, firm's incentives are aligned with the planner in the good state while in the bad state $v_f > v_p^B$. Therefore, as shown in Lemma 4, f's lower best reply $b_f(S)$ is to the left of $b_p(S)$, thus implying that the firm has excessive incentives to experiment, $s_{pp} < s^*$.

- In the other polar case, if $e^G = v_p^G$ (or equivalently $v_f = 0$), firm f has no incentive to experiment ex ante and thus $s_{pp} = S^* > s^*$. By continuity, there exists $\hat{e} \in (0, v_p^G)$ such that if $e^G = \hat{e}$ then $s_{pp} = s^*$. Moreover, $\hat{e}$ is unique since $\frac{\partial b_f(S)}{\partial v_f} < 0$ by Lemma 3.

(a) The preliminary results above show that, if $e^G < \hat{e}$ and regulatory tools are chosen at their first-best levels, the firm has excessive experimentation incentives $s_{pp} < s^*$. A marginal increase in S_{pp} above S^* and a marginal increase of z_{pp} above z^* induce second-order losses, which are trumped by first-order gains:

- By Lemma 6 $b_f(S)$ is increasing in z if and only if $z > 0$.[48] Setting $z_{pp} > z^*$ results in a first-order gain by moving s_{pp} closer to s^*.
- By Lemma 2 the lower best reply is increasing in S for $S > S^*$.[49] Setting $S_{pp} > S^*$ results in a first-order gain by moving s_{pp} closer to s^*.

[48]Note that for player f the optimal withdrawal is zero i.e., never withdraw.

[49]Under withdrawal regulation f's incentives in the bad state are higher $v_f > v_p^B$, therefore from Lemma 4, we know that $B_f(s)$ lies below the planner one $B_p(s)$. Hence, the optimal approval that f would choose is lower than S^*.

(b) If $e^G > \hat{e}$ firm f's experimentation incentives are insufficient $s^* < s_{pp}$. Applying the same arguments as in point (a), we can show that the approval and withdrawal standards should be set at $S_{pp} < S^*$ and $z_{pp} < z^*$ to encourage experimentation.

Proof of Proposition 5
(a) At $e^G = 0$ the optimal liability $L^* = \overline{L}$ achieves the first best since p and f incentives in the good state are perfectly aligned. Authorization regulation, instead, does not achieve the first best. As shown in Proposition 4, the firm has excessive incentives to experiment in the first phase. At $e^G = \hat{e}$, by Proposition 4 authorization regulation achieves the first best, while optimal liability leads to an insufficient level of experimentation. By continuity, there exists an intermediate value $\tilde{e}$ such that the planner uses liability if $e^G < \tilde{e}$.

(b) - (c) The welfare achieved when the planner has access to all three instruments is represented in Figure 5. For $e^G = 0$, the first best is achieved setting $L_m^* = \overline{L}$. For $e^G = \hat{e}$, we showed that the first best is achieved using only authorization regulation with $S_m^* = S^*$, $z_m^* = z^*$, and zero liability $L_m = 0$.

Consider now $e^G \in (0, \hat{e})$ and suppose the planner commits to $S_m = S^*$ and $z_m = z^*$. If $L_m = 0$, we know from Proposition 4 that f's experimentation incentives are excessive $s_m < s^*$. If instead $L_m = \overline{L}$, using Lemma 3 and Lemma 4, experimentation incentives are insufficient $s^* < s_m$. By continuity, there exists $L_m^* \in (0, \overline{L})$ such that $s_m = s^*$. This L_m^* is also unique since L_m is decreasing in e^G. Overall, the planner achieves the first best by committing to $S_m^* = S^*$, $s_m^* = s^*$, and L_m^*.

At $e^G \in (\hat{e}, v_p^G)$, f's incentives are lower in both states. If $S_m = S^*$ and $z_m = z^*$, using Lemma 3 and Lemma 4, we know f's experimentation incentives are insufficient $s^* < s_m$. Liability has to be preempted $L_m^* = 0$ since any positive liability would further reduce f's experimentation incentives. Finally, using results of Proposition 4, we have that for $e^G \in (\hat{e}, v_p^G)$ optimal authorization regulation has to be lenient, $S_m < S^*$ and $z_m < z^*$.

C Appendix C: Proofs for Appendix A

Proof of Lemma 1

For the first part of the result about the ex ante phase see the proof of Lemma B0 in the Online Appendix of Henry and Ottaviani (2019). The second part of the result about the ex post phase follows similar steps.

Proof of Lemma 2
(a) We distinguish three cases:

(i) Let $S \in [0, z^*]$ where z^* is the belief at which $u_j^{\mathrm{II}}(z^*)$ is tangent at 0. By way of contradiction suppose $s = b_j(S) < S$. Since $S < z^*$, we have $u_j^{\mathrm{II}}(S) = 0$ and this implies that experimenting for $q \in (s, S)$ generates a cost without any informational gain. Player j would be better off immediately rejecting at any $q \in (s,S)$ which guarantees a payoff of zero. Hence, it must be that $s = b_j(S) = S$.

(ii) Let $S > z^*$. The ex post value function $u_j^{\mathrm{II}}(q)$ (the blue dot-dashed curve in Figure 9) is constant at 0 for $q \in [0,z^*]$, increasing and positive for $q \in (z^*,1]$. Therefore, $u^{\mathrm{II}}(S) > 0$ and information is now valuable. The solution $s = b_j(S)$ is interior and solves (12). The lower best reply function (the red dashed curve in Figure 9) is first decreasing, reaches a minimum at $s = s^*$, when $S = S^*$, and then increases to reach 1 in the limit as $S \to 1$. Given S, increasing the rejection standard s reduces experimentation costs but increases the probability that the rejection standard is reached before the approval standard with an associated payoff of zero instead of $u_j^{\mathrm{II}}(S) > 0$. When $S = z^*$ from part (i) we know that $s = b_j(z^*) = z^*$. As S increases toward S^*, $b_j(S)$ decreases toward s^* to get closer to the first-best value function (solid blue line) i.e., adoption and experimentation are strategic substitute. Figure 9 shows that, when constrained to take $\underline{S} \in (z^*,S^*)$, player j best replies with $s = b_j(\underline{S})$. The associated ex ante value function, plotted in thick-yellow, lies everywhere below the first best ex ante value function in solid blue, is tangent to the 0 line at $s = b_j(\underline{S})$ and intersects the ex post value function with a kink at $\underline{S}$. As $\underline{S} \to S^*$, $b_j(S) \to s^*$ and the yellow ex ante value function converges to the blue first-best ex ante value. At $S = S^*$, player j best replies $b_j(S^*) = s^*$ achieving the first-best solution.

(iii) As S increases from S^* to 1, $b_j(S)$ increases from s^* to 1. In this case, approval and experimentation are strategic complements. As shown in Figure 9, when constrained to take $\overline{S} \in (s^*,1]$, player j best replies with $s = b_j(\overline{S})$. The associated ex ante value function, plotted in green, lies everywhere below the first-best ex ante value function, is tangent to the 0 line at s and intersects the ex post value function with a kink at $\overline{S}$. Note that $s = b_j(\overline{S}) = b_j(\underline{S})$ is the best reply to both $\underline{S}$ and $\overline{S}$. Thus, the lower best reply is hump shaped, with the yellow and green ex ante values overlapping for $q \in (s,\underline{S})$. As $\overline{S} \to 1$, $b_j(S) \to 1$ in order to reduce the experimentation cost, in the limit when $\overline{S} = 1$, $b_j(1) = 1$.

(b) The proof proceed by cases:

(i) For $s \in [0,s^*)$ player j can always guarantee himself the solid blue ex-ante value in figure 10 by approving and immediately withdrawing for $q \leq s^*$, experimenting for $q \in (s^*,S^*)$ and approving for $q \geq S^*$. In Figure 10 we plot the upper best reply with the convention that q lies to the right of s^* for any s. With this convention $B_j(s) = S^*$, player j approves as soon as q hits S^* or approves and withdraws as soon as q hits s^*. Overall, the constraint of rejecting at s does not bind for $s < s^*$. This is because a player in charge of approval only is nonetheless able to control rejections by approving and suddenly withdrawing. Thus, $B_j(s) \equiv S^*$ whenever $s \in [0,s^*)$.

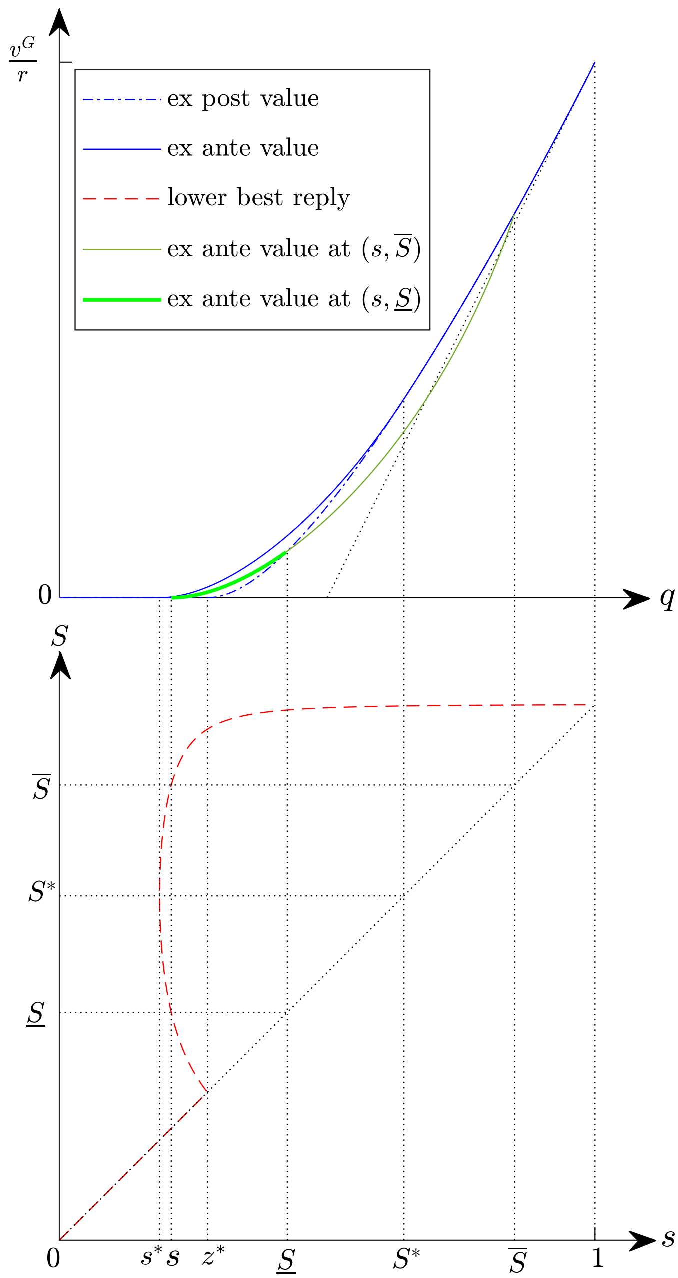

Figure 9: Construction of the lower best reply (bottom panel) in relation to the corresponding value functions (top panel).

(ii) As s increases from s^* to $\bar{s}$, $B_j(s)$ decreases toward the diagonal. At $\hat{s} \in (s^*, \bar{s})$ the ex ante value in dark-yellow lies below the socially optimal value, is tangent to the ex post value at $S = B_j(\hat{s})$ and intersects the 0 line at $\hat{s}$. As s increases from $\hat{s}$ the yellow value decreases toward the blue dashed-dotted ex post value, at $\bar{s}$ there is no non-empty interval of belief for which the yellow value is above the ex post value. For $s \in (z^*, \bar{s})$ we still have that $B_j(s) > s$. Here player j finds optimal to approves at both $q \geq B_j(s)$ and $q \leq s$ while experimenting for $q \in (s, B_j(s))$. This allows player j to obtain $u_j^{\mathrm{I}}(s) > 0$ instead of obtaining 0 following immediate rejection.[50]

(iv) For $s \in (\bar{s}, 1]$ the best reply lies on diagonal, $B_j(s) = s$ since there is no non-empty interval of beliefs for which the ex ante value lies above the ex post value. In other words, there is no belief for which ex ante experimentation is worthwhile.

Proof of Lemma 3

(a) The fact z^* decreases with v_j^G follows directly from expression (11). An increase in v_j^G, increases the payoffs of player j on the market and delays withdrawal.

(b) For $\mathsf{S} > \mathsf{z}^*$, $b_j(\mathsf{S})$ is the value of s that solves $\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{s}} = 0$. By the implicit function theorem

$$\frac{\partial b_j(\mathsf{S})}{\partial v_j^G} = -\frac{\frac{\partial^2 u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{s} \partial v_j^G}}{\frac{\partial^2 u_j^{\mathrm{I}}(\sigma)}{\partial^2 \mathsf{s}}}.$$

The concavity of the problem[51] implies that the sign of $\frac{\partial b_j(\mathsf{S})}{\partial v_j^G}$ is determined by the sign of $\frac{\partial^2 u_j^{\mathrm{I}}}{\partial \mathsf{s} \partial v_j^G}$ which is given, using expression (9) for u_j^{I}, by

$$\frac{\partial^2 u_j^{\mathrm{I}}}{\partial \mathsf{s} \partial v_j^G} = \frac{\partial \Psi^{\mathrm{I}}}{\partial \mathsf{s}} \frac{\partial u_j^{\mathrm{II}}(\mathsf{S})}{\partial v_j^G}. \tag{16}$$

Next, we show that $\frac{\partial^2 u_j^{\mathrm{I}}}{\partial s \partial v_j^G} < 0$; increasing v_j^G raises the value of experimentation without affecting expected costs. According to Lemma 1 the first term of expression (16) $\frac{\partial \Psi^{\mathrm{I}}}{\partial s}$ is negative; for given σ and S, increasing s reduces the probability that S is reached first. The second factor $\frac{\partial u_j^{\mathrm{II}}}{\partial v_j^G} = \frac{e^{\mathsf{S}}}{1+e^{\mathsf{S}}} \frac{1}{r} \left[1 - \psi^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*)\right] > 0$ is positive because the ex post value u_j^{II} is increasing in v_j^G. Overall, the lower best reply shifts to the left as v_j^G increases from $v_{j,l}^G$ to $v_{j,h}^G$, as shown in Figure 6.

[50] In this case $B_j(s)$ solves (13) with the additional term $g\left(u_j^{\mathrm{II}}(\mathsf{s}, G, \mathsf{z}^*) + e^{-\mathsf{s}} u_j^{\mathrm{II}}(\mathsf{s}, B, \mathsf{z}^*)\right)$.

[51] See Henry and Ottaviani (2019).

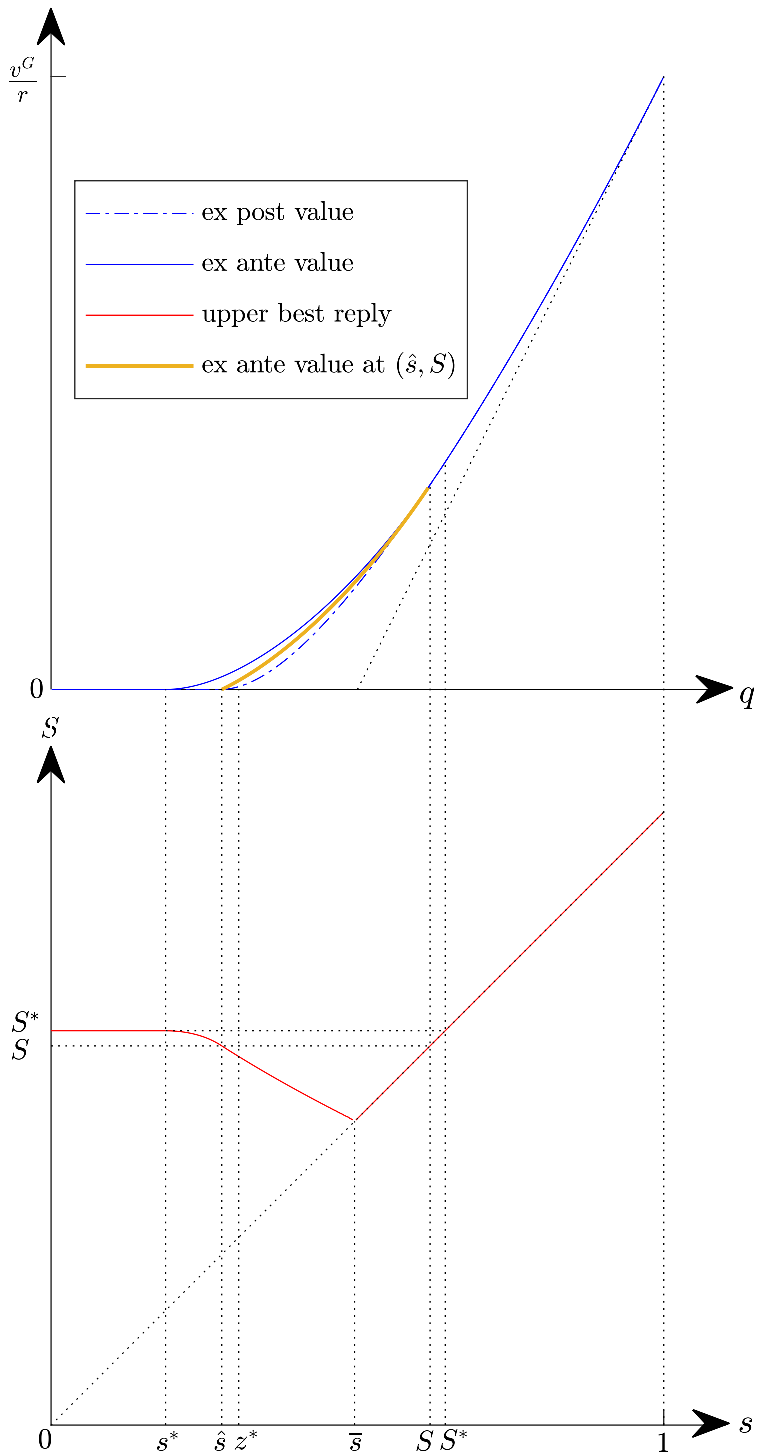

Figure 10: Construction of upper best reply (bottom panel) in connection to value functions (top panel).

(c) As in (b), by the implicit function theorem, the sign of $\frac{\partial B(\mathsf{s})}{\partial v^G}$ is determined by the sign of $\frac{\partial u_j^{\mathrm{I}}}{\partial \mathsf{S} \partial v^G}$. Since z^* is a function of v_j^G, this can be decomposed in two effects 1) direct effect 2) indirect effect through z^*

$$\underbrace{\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{S} \partial v^G}}_{\text{direct effect}} + \underbrace{\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{S} \partial \mathsf{z}} \frac{\partial \mathsf{z}^*}{\partial v^G}}_{\text{effect through } \mathsf{z}^*}. \tag{17}$$

Direct Effect. The direct effect can be decomposed in two parts. First, an increase in v_j^G increases the opportunity cost of delaying approval. There is an effect going in the opposite direction: the value of information increases thus encouraging to delay approval. We show below the first term dominates. At $\mathsf{z} = \mathsf{z}^*$, taking derivatives of (13) with respect to v_j^G we have $\frac{\partial^2 u_j^{\mathrm{I}}}{\partial \mathsf{S} \partial v_j^G} = \frac{e^\sigma}{1+e^\sigma} \Psi^{\mathrm{I}}(\sigma, G) \frac{f}{r} \left[1 - \psi^{\mathrm{II}}(\mathsf{S}, z^*, G)\right] < 0$, which is negative by $f < 0$. This proves that the direct effect is negative.

Indirect Effect. We now show that the indirect effect is zero for $\mathsf{z} = \mathsf{z}^*$. Taking derivatives of equation (9) with respect to z yields

$$\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{z}} = \Psi^{\mathrm{I}}(\sigma) \frac{\partial u_j^{\mathrm{II}}(\mathsf{S})}{\partial \mathsf{z}}$$

and thus

$$\frac{\partial^2 u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{S} \partial \mathsf{z}} = \frac{\partial \Psi^{\mathrm{I}}(\sigma)}{\partial \mathsf{S}} \frac{\partial u_j^{\mathrm{II}}(\mathsf{S})}{\partial \mathsf{z}} + \Psi^{\mathrm{I}}(\sigma) \frac{\partial^2 u_j^{\mathrm{II}}(\mathsf{S})}{\partial \mathsf{S} \partial \mathsf{z}}. \tag{18}$$

The first term is zero for $\mathsf{z} = \mathsf{z}^*$ since, by definition, z^* is the value of z that maximizes u_j^{II}. Furthermore $u_j^{\mathrm{II}}(\mathsf{S})$ is proportional to

$$u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}) + e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}),$$

so that

$$\frac{\partial^2 u_j^{\mathrm{II}}(\mathsf{S})}{\partial \mathsf{S} \partial \mathsf{z}} = \left[\frac{\partial^2 u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{S} \partial \mathsf{z}} + e^{-\mathsf{S}} \frac{\partial^2 u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{S} \partial \mathsf{z}}\right] - e^{-\mathsf{S}} \frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}}. \tag{19}$$

Using the results of Lemma 1 characterizing hitting probabilities, we can show that $\frac{\partial^2 u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{S} \partial \mathsf{z}} = -r_2 \frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{z}}$ and $\frac{\partial^2 u^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{S} \partial \mathsf{z}} = r_1 \frac{\partial u^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}}$. So overall, we can rewrite (19), as

$$\begin{aligned}
\frac{\partial^2 u_j^{\mathrm{II}}(\mathsf{S})}{\partial \mathsf{S} \partial \mathsf{z}} &= -e^{-\mathsf{S}} \frac{\partial u^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}} + \left[-r_2 \frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{z}} + e^{-\mathsf{S}} r_1 \frac{\partial u^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}}\right] \\
&= -r_2 \left[\frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{z}} + e^{-\mathsf{S}} \frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}}\right] \qquad (20)
\end{aligned}$$

where the last equality follows from the fact $r_1 = 1 - r_2$. By definition of z^*, equation (20) is thus 0 at $\mathsf{z} = \mathsf{z}^*$.

Overall this implies that the indirect effect is zero. Since we showed that the direct effect is negative, we conclude that the upper best reply shifts down as v_j^G increases from $v_{j,l}^G$ to $v_{j,h}^G$, as shown in Figure 7.

Proof of Lemma 4

(a) The fact that z^* decreases follows directly from expression (11) for z^*. An increase in v_j^B increase the expected payoff while on the market at any belief and delays withdrawal.

(b) For $\mathsf{S} > \bar{\mathsf{z}}$, $b_j(\mathsf{S})$ is the value of s that solves $\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{s}}$. As before, by the implicit function theorem and the concavity of the problem the sign of $\frac{\partial b_j(\mathsf{S})}{\partial v_j^B}$ is determined by the sign of $\frac{\partial^2 u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{s} \partial v_j^B}$ which is given, using expression (9), by

$$\frac{\partial^2 u_j^{\mathrm{I}}}{\partial v_j^B \partial \mathsf{s}} = \frac{\partial \Psi^{\mathrm{I}}}{\partial \mathsf{s}} \frac{\partial u_j^{\mathrm{II}}(\mathsf{S})}{\partial v_j^B}. \tag{21}$$

According to Lemma (1) the first term of expression (21) $\frac{\partial \Psi^{\mathrm{I}}}{\partial \mathsf{s}}$ is negative. Intuitively, given σ and S, increasing s reduces the probability that S is reached first. The second term $\frac{\partial u_j^{\mathrm{II}}(\mathsf{S})}{\partial v_j^B} = \frac{1}{1+e^{\mathsf{S}}} \frac{1}{r}(1 - \psi^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*)) > 0$ is positive since the ex post value u_j^{II} is increasing in v_j^B.

(c) As in (b), by the implicit function theorem, the sign of $\frac{\partial B(s)}{\partial v_j^B}$ is determined by the sign of $\frac{\partial u_j^I}{\partial \mathsf{S} \partial v_j^B}$. Since z^* is a function of v_j^B, this can be decomposed in two effects 1) direct effect 2) indirect effect through z^*

$$\underbrace{\frac{\partial u_j^I(\sigma)}{\partial \mathsf{S} \partial v_j^B}}_{\text{direct effect}} + \underbrace{\frac{\partial u_j^I(\sigma)}{\partial \mathsf{S} \partial \mathsf{z}} \frac{\partial \mathsf{z}}{\partial v_j^B}}_{\text{effect through } \mathsf{z}^*}. \tag{22}$$

Direct Effect. The direct effect can be decomposed in two parts. First, an increase in v_j^B increases the opportunity cost of delaying approval. There is an effect going in the opposite direction: the value of information increases thus encouraging to delay approval. We show below that the first term dominates. Following the same reasoning in Lemma (3) taking derivatives of (13) with respect to v_j^B, $\frac{\partial^2 u_j^{\mathrm{I}}}{\partial \mathsf{S} \partial v_j^B} = \frac{e^{\sigma}}{1+e^{\sigma}} \Psi^{\mathrm{I}}(\sigma, G) \frac{(f-1)e^{-\mathsf{S}}}{r} \left[1 - \psi^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*)\right] < 0$ which is negative since $f < 0$. This proves that the direct effect is negative.

Indirect Effect. Following steps analogous to Lemma (3), to show that the indirect effect is zero it is enough to focus on the term $\frac{\partial^2 u_j^{\mathrm{II}}}{\partial \mathsf{S} \partial \mathsf{z}}$ which can be rewritten as

$$\frac{\partial^2 u_j^{\mathrm{II}}}{\partial \mathsf{S} \partial \mathsf{z}} = -r_2 \left[\frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{z}} + e^{-\mathsf{S}} \frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}} \right]. \tag{23}$$

At $\mathsf{z} = \mathsf{z}^*$ the term in the brackets is zero, thus implying that the indirect effect is zero. Overall, since we showed that the direct effect is negative, the upper best reply shifts down as v_j^B increases.

Proof of Lemma 5

We derive comparative statics with respect to α. By the implicit function theorem and the concavity of the problem the sign of $\frac{\partial b_j(\mathsf{S})}{\partial \alpha}$ is determined by $\frac{\partial^2 u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{s} \partial \alpha}$. Taking the derivative of (12) with respect to α, since $\frac{\partial u_j^{\mathrm{II}}}{\partial \alpha} = \frac{1}{\alpha} u_j^{\mathrm{II}}$:

$$\frac{\partial^2 u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{s} \partial \alpha} = \frac{e^{\sigma}}{1+e^{\sigma}} \psi(\sigma, G) \frac{a}{\alpha} \left[u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) \right] < 0$$

which is negative because $a < 0$.

Similarly, by the implicit function theorem and the concavity of the problem the sign of $\frac{\partial B_j(\mathsf{s})}{\partial \alpha}$ is determined by

$$\frac{\partial^2 u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{S} \partial \alpha} = \frac{e^{\sigma}}{1+e^{\sigma}} \Psi(\sigma, G) \frac{f}{\alpha} \left[u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) - e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) \right] > 0 \tag{24}$$

which is positive because $f < 0$ and the term $u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) - e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*)$ is negative as we show next.

At $\mathsf{S} = B_j(\mathsf{s})$ equation (13) is satisfied. This implies

$$f \left[u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) - e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) \right] = - \left[f(1+e^{-\mathsf{S}}) \frac{c}{r} - e^{-\mathsf{S}} \frac{c}{r} + g(1+e^{-\mathsf{s}}) \frac{c}{r} \right]. \tag{25}$$

The term in brackets on the right hand side of (25) is proportional to $\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{S}}$ for a player j with $v_j^G = v_j^B = 0$. Clearly, for this player experimentation is not valuable and it must be that $B_j(\mathsf{s}) = \mathsf{s}$ or equivalently $\frac{\partial u_j^{\mathrm{I}}(\sigma)}{\partial \mathsf{S}} < 0$. Overall, the right hand side of (25) is positive thus implying that the term $u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z}^*) + e^{-\mathsf{S}} u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z}^*) - e^{-\mathsf{S}}(\mathsf{S}, B, \mathsf{z}^*)$ is negative.

Proof of Lemma 6

For $\mathsf{S} > \mathsf{z}$, $b_j(\mathsf{S})$ is the value of s that solves $\frac{\partial u_j^{\mathrm{I}}}{\partial \mathsf{s}} = 0$. By the implicit function theorem and the concavity of the problem the sign of $\frac{\partial b_j(\mathsf{S})}{\partial \mathsf{z}}$ is determined by $\frac{\partial^2 u_j^{\mathrm{I}}}{\partial \mathsf{s} \partial \mathsf{z}}$. Taking the cross-derivative of (9) with respect to z and s we obtain

$$\frac{\partial^2 u_j^{\mathrm{I}}}{\partial \mathsf{s} \partial \mathsf{z}} = \frac{\partial \Psi^{\mathrm{I}}}{\partial \mathsf{s}} \frac{\partial u_j^{\mathrm{II}}(\mathsf{S})}{\partial \mathsf{z}}. \tag{26}$$

From Lemma (1) the first term $\frac{\partial \Psi^{\mathrm{I}}}{\partial \mathsf{s}} < 0$ is negative. Intuitively, given σ and S increasing s reduces the probability that S is reached first. The second term $\frac{\partial u_j^{\mathrm{II}}(\mathsf{S})}{\partial \mathsf{z}}$, being the first order condition of the ex post problem, is positive whenever $\mathsf{z} < \mathsf{z}^*$. Overall, (26) is positive if and only if $\mathsf{z} > \mathsf{z}^*$.

Turning to the upper best reply for $\mathsf{s} < \tilde{\mathsf{s}}$, $B_j(\mathsf{s})$ is the value of S that solves $\frac{\partial u_j^{\mathrm{I}}}{\partial \mathsf{S}} = 0$. By the implicit function theorem and the concavity of the problem the sign of $\frac{\partial B_j(\mathsf{z})}{\partial \mathsf{z}}$ is determined by

$\frac{\partial u_j^{\mathrm{I}}}{\partial \mathsf{S} \partial \mathsf{z}}$. Taking the derivative of (13) with respect to z we obtain

$$\frac{\partial^2 u_j^{\mathrm{I}}}{\partial \mathsf{S} \partial \mathsf{z}} = \frac{e^{\sigma}}{1+e^{\sigma}} \Psi^{\mathrm{I}}(\mathsf{S}, G) \Big[f \left[\frac{u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{z}} + e^{-\mathsf{S}} \frac{u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}} \right] - e^{-\mathsf{S}} \frac{u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}} \tag{27}$$

$$+ \left[\frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{S} \partial \mathsf{z}} + e^{-\mathsf{S}} \frac{\partial u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{S} \partial \mathsf{z}} \right] \Big], \tag{28}$$

Using the results of Lemma 1 characterizing hitting probabilities, we can show that $\frac{\partial^2 u_j^{\mathrm{II}}(\mathsf{S},G,\mathsf{z})}{\partial \mathsf{S} \partial \mathsf{z}} = -r_2 \frac{\partial u_j^{\mathrm{II}}(\mathsf{S},G,\mathsf{z})}{\partial \mathsf{z}}$ and $\frac{\partial^2 u^{\mathrm{II}}(\mathsf{S},B,\mathsf{z})}{\partial \mathsf{S} \partial \mathsf{z}} = r_1 \frac{\partial u^{II}(\mathsf{S},B,\mathsf{z})}{\partial \mathsf{z}}$. This together with the fact that $r_1 + r_2 = 1$ we have

$$\frac{\partial^2 u_j^{\mathrm{I}}}{\partial \mathsf{S} \partial \mathsf{z}} = \frac{e^{\sigma}}{1+e^{\sigma}} \Psi^{\mathrm{I}}(\mathsf{S}, G)(f - r_2) \left[\frac{u_j^{\mathrm{II}}(\mathsf{S}, G, \mathsf{z})}{\partial \mathsf{z}} + e^{-\mathsf{S}} \frac{u_j^{\mathrm{II}}(\mathsf{S}, B, \mathsf{z})}{\partial \mathsf{z}} \right]. \tag{29}$$

The term in brackets is negative whenever $\mathsf{z} > \mathsf{z}^*$ being the first order condition of the ex post problem. Overall, since $f < 0$, the expression above is positive whenever $\mathsf{z} > \mathsf{z}^*$.

D Appendix D: Model with Adverse Events

This appendix formulates a model for the arrival of information in the second stage alternative to the Wiener process considered in the paper. Suppose instead that in the second stage bad news arrives according to a Poisson process. Specifically, assume that in state B the planner obtains a negative payoff $-D$ with arrival time distributed according to an exponential distribution with parameter λ.

In this environment, learning in the second phase is different from the baseline. On the one hand, when no adverse event occurs, players become increasingly confident that the state is G. On the other hand, an adverse event perfectly reveals that the state is B and triggers immediate withdrawal by the planner. The incentives of the firm may be (partially) aligned by imposing liability L as soon as an adverse event occurs.

Once the adverse event occurs, the firm incurs liability but has no incentive to subsequently withdraw from the market. Withdrawal has to be immediately ordered by the planner following the adverse event. Before the occurrence of the adverse event, the firm will never withdraw, as it becomes increasingly confident that the state is G. The planner is thus effectively constrained to using liability and approval regulation.[52]

The first phase experimentation problem becomes a standard Wald problem in which upon approval the planner obtains payoff $v_p^G = \frac{v_p}{r} > 0$ in the good state and $v_p^B = \frac{v_p}{r} - \frac{\lambda}{\lambda+r}(v_p + D)$ in the bad state, where D is large enough so that $v_p^B < 0$. The firm obtains $v_f^G = \frac{v_f}{r} > 0$ and

[52] Learning in this model is too coarse to allow us to analyze the ex ante implications of the stringency of withdrawal regulation.

$v_f^B = \frac{v_f}{r} - \frac{\lambda}{\lambda+r}(v_f + L)$. The model can be seen as a special case of the model analyzed in the main text, where the withdrawal dimension is removed. The same intuition underlying Proposition 5 applies. When e^G is small, given an approval threshold chosen at the fist best level $S_{pp} = S^*$, the liability can be adjusted to induce the firm to experiment up to threshold $s_{pp} = s^*$. However, when e^G is larger, liability needs to be preempted to provide experimentation incentives to the firm. As an immediate corollary of Proposition 5, we obtain:

Proposition 6 **(a)** ***Externality Criterion:*** *When the planner has to choose between liability and authorization regulation, there exists $\tilde{e}_2$ such that the planner uses liability if $e^G < \tilde{e}_2$.*

(b) ***Preemption:*** *When the planner has access to both tools, there exists $\hat{e}_2$ such that the socially optimal mix (S_m^*, L_m^*) is such that liability is preempted ($L_m = 0$) if and only if $e^G \geq \hat{e}_2$. For $e^G < \hat{e}_2$, liability is not preempted ($L_m^* > 0$) and decreasing in e^G.*

(c) ***Lenient Regulation Property:*** *for $e^G > \hat{e}_2$ approval regulation is more lenient compared to the first best ($S_{pp}^* < S^*$).*

Proof of Proposition 6

In this setting, under authorization regulation, the planner controls the approval standard and the firm chooses the experimentation standard. Withdrawal is not relevant because (i) once the adverse event occurs, the state is learned and withdrawal occurs and (ii) if no adverse event has occurred, players become increasingly confident that the state is G.

Suppose approval is chosen at the first-best level $S_{pp} = S^*$:

- If $e^G = 0$, firm's incentives are aligned with the planner in the good state while in the bad state $v_f > v_p^B$. Therefore, as shown in Lemma 4, f's lower best reply $b_f(S)$ is to the left of $b_p(S)$, thus implying that the firm has excessive incentives to experiment, $s_{pp} < s^*$.

- In the other polar case, if $e^G = v_p^G$ (or equivalently $v_f = 0$), firm f has no incentive to experiment ex ante and thus $s_{pp} = S^* > s^*$. By continuity, there exists $\hat{e}_2 \in (0, v_p^G)$ such that if $e^G = \hat{e}_2$ then $s_{pp} = s^*$.

- Thus, if $e^G < \hat{e}_2$ and $S_{pp} = S^*$, the firm has excessive experimentation incentives $s_{pp} < s^*$. A marginal increase in S_{pp} above S^* induces second-order losses, which are trumped by first-order gains. By Lemma 2 the lower best reply is increasing in S for $S > S^*$. Setting $S_{pp} > S^*$ results in a first-order gain by moving s_{pp} closer to s^*. If $e^G > \hat{e}_2$ firm f's experimentation incentives are insufficient $s^* < s_{pp}$. Applying the same arguments, the approval standard should be set at $S_{pp} < S^*$ to encourage experimentation.

With these preliminary results we now prove parts (a)-(c).

(a) At $e^G = 0$ the optimal liability $L^* = \overline{L}$ achieves the first best since p and f incentives in the good state are perfectly aligned. Authorization regulation, instead, does not achieve the first best. At $e^G = \hat{e}_2$, by the preliminary result above, authorization regulation achieves the first best, while optimal liability leads to an insufficient level of experimentation. By continuity, there exists an intermediate value $\tilde{e}_2$ such that the planner uses liability if $e^G < \tilde{e}_2$.

(b) - (c) For $e^G = 0$, the first best is achieved setting $L_m^* = \overline{L}$. For $e^G = \hat{e}_2$, we showed that the first best is achieved using only authorization regulation with $S_m^* = S^*$ and zero liability $L_m = 0$.

Consider now $e^G \in (0, \hat{e}_2)$ and suppose the planner commits to $S_m = S^*$. If $L_m = 0$, we know from the preliminary results that f's experimentation incentives are excessive $s_m < s^*$. If instead $L_m = \overline{L}$, using Lemma 3 and Lemma 4, experimentation incentives are insufficient $s^* < s_m$. By continuity, there exists $L_m^* \in (0, \overline{L})$ such that $s_m = s^*$. This L_m^* is also unique since L_m is decreasing in e^G. Overall, the planner achieves the first best by committing to $S_m^* = S^*$ and L_m^*.

At $e^G \in (\hat{e}_2, v_p^G)$, f's incentives are lower in both states. If $S_m = S^*$, using Lemma 3 and Lemma 4, we know f's experimentation incentives are insufficient $s^* < s_m$. Liability has to be preempted $L_m^* = 0$ since any positive liability would further reduce f's experimentation incentives. Finally, using the preliminary results above, we have that for $e^G \in (\hat{e}_2, v_p^G)$ optimal authorization regulation has to be lenient, $S_m < S^*$.