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

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

November 2018

Abstract

Dynamic adoption policies of activities with uncertain returns are characterized by three key decisions: in the ex ante experimentation phase, the decisions when to abandon experimentation and when to introduce to market; in the ex post learning phase, the decision when to withdraw following the accumulation of bad news. In a tractable continuous-time model, we study the optimal mix of the three instruments regulators employ to align the private incentives of firms: ex ante approval regulation, ex post withdrawal regulation, and liability. Our results can rationalize the array of regulatory environments observed across applications ranging from product safety to patent protection. We also consider costly lying and show that the social planner can be better off when the firm privately observes research results.

Keywords: Experimentation, approval regulation, liability, withdrawal.

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

[*]We thank Daniel Carpenter, Andrew Daughety, Xinyu Hua, Kathryn Spier, and seminar participants at EALE 2018 in Milan, ESSET 2017 in Gerzensee, UCL, and Ecole Polytechnique 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

In *Wyeth v. Levine*, 555 U.S. 555 (2009), the US Supreme Court held 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 so-called preemption doctrine argued that exposing pharmaceutical companies to liability would reduce innovation incentives prior to market introduction, while opponents maintained that the threat of liability would induce more timely information disclosure and voluntary withdrawals of harmful drugs; see Garber (2013). The controversy surrounding this ruling illustrates the difficulty of striking a balance between ex ante approval regulation, ex post withdrawal regulation, and liability.[1]

To analyze the law and economics of regulation of new drugs and other activities with uncertain returns, this paper formulates a novel two-phase model with ex ante experimentation and ex post learning. First, a phase of preliminary experimentation with market research and testing takes place prior to market introduction. This initial experimentation is carried out in a rather controlled environment; think of carefully monitored clinical trials.[2] In the ex ante experimentation phase, the decision is eventually taken to either abandon experimentation or to introduce to market. Second, in the ex post learning phase after market introduction, learning continues, often in a less controlled way and at a different speed, possibly leading to a final withdrawal decision following the accumulation of sufficiently bad news.[3]

Our main specification features a planner who regulates a firm undertaking an activity that generates an externality of uncertain sign. We capture the uncertainty in the sign of the externality by positing that the firm does not suffer the full social damage in the bad state and does not recoup the full social benefits in the good state. The role of regulation is then to re-align misaligned private incentives with social incentives. Our key focus is to determine the optimal mix of the three main regulatory instruments that are typically used in practice to regulate drugs and other activities with uncertain externalities: liability, approval (ex ante) regulation, and withdrawal (ex post) regulation.

The mix of these regulatory instruments vary across applications, as we argue in the institutional analysis in Section 4. Three main regularities emerge from the summary in Table 1. First, when regulation is used either ex ante or ex post, it tends to be lenient, as with post-

[1] The preemption doctrine, also known as "regulatory compliance defence," prevails not only for medical devices, following *Riegel v. Medtronic*, 552 U.S. 312 (2008), but also for generic drugs. Furthermore, even for regular drugs, application of preemption varies by states. Garber (2013) reports that "several state high courts such as those of California, Washington and Utah have adopted, as a practical matter, a regulatory compliance defense for prescription drugs."

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

[3] In a similar vein, learning in this second phase could result also from "wet bench" experiments or observational learning.

	Ex Ante Regulation	Ex Post Regulation	Liability
Pharmaceutical Safety	clinical trials	surveillance (limited)	some
Product Safety	little	largely voluntary	mostly
Patent validity	patent office (lenient)	courts or patent office (lenient)	no
Zoning/Environment	mostly	limited/irreversible	some

Table 1: Regulatory environments across applications.

marketing pharmacovigilance or patent grants and patent invalidations. Second, liability and ex post regulation are rarely used in conjunction. Third, liability tends to be common in areas where externalities are small (e.g., introduction of regular products), while regulation dominates with large externalities (e.g., drugs, patents, and environmental projects).

The continuous-time model we formulate in Section 2 captures in a stylized and tractable way the main features of the regulatory environments that prevail in these applications. Firm i collects information, which is publicly observed, in the form of a Wiener process whose drift depends on a binary state of the world, either good or bad. Information acquisition is costly in the first phase and costless after implementation. Decision payoffs are collected only 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 $v_p^B < 0$, while the firm, before any liability is imposed, obtains a positive payoff $v_i > 0$ that is independent of the state but does not capture the full social benefits in the good state, $v_i < v_p^G$. For the purpose of our main comparative statics exercise, we set $v_i = (1-e)v_p^G$ so that e parametrizes the fraction of the planner's payoff not internalized by the firm (i.e., the extent of the positive externality) in the good state.

Section 3 sets the stage by characterizing the planner solution, the first-best welfare benchmark for the model. Even though information is costless after adoption, the planner uses a balance between ex ante testing and ex post monitoring, given that information is indirectly costly in the ex post phase in the bad state. 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. In the ex ante phase, the model corresponds to a Wald problem, characterized by two cutoffs, a research cutoff s, such that the planner abandons and rejects the project as soon as the belief q falls below s, and an adoption cutoff S such that the activity is implemented as soon as the belief reaches S. The social planner collects information for beliefs between these two cutoffs $q \in (s, S)$, which are centered around the withdrawal cutoff $z \in (s, S)$. The planner strikes a balance between ex ante testing and ex post surveillance.

We then turn to studying the performance of different regulatory instruments in aligning the incentives of the firm with those of the planner and use our model to explain the regularities observed in practice. We first examine in Section 5 the properties of the instruments considered individually. Consider first liability, understood as a liability rate charged for each unit of time in the ex post phase (in which a product is sold on the market for instance) when the state is bad, so that the flow payoff in that state becomes $v_i - L$. Without externality ($e = 0$), the first best is clearly achieved by a strong liability rule requiring the firm to fully compensate for the whole social damage in the bad state, because the incentives are then perfectly aligned in both states. However, in the presence of a positive externality ($e > 0$), it is optimal for the planner to commit to being more lenient than under strong liability, in order not to discourage research.

By appropriately choosing 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 the socially optimal withdrawal z^*. We show that if $e > 0$, it is optimal for the planner to deviate from $\hat{L}$. Indeed, when liability is set at $\hat{L}$, the firm rejects too early but also adopts too early in the ex ante phase. Increasing or decreasing liability thus has conflicting effects on research and approval. If the initial belief q_0 is low, it is optimal for the planner to be more lenient since encouraging experimentation is the most pressing issue. If q_0 is high, it is optimal for the planner to be tougher so as to discourage early approval.

We show that liability can be chosen to induce any withdrawal decision and is thus a substitute instrument to ex post regulation. However, the two tools have very different implications for ex ante incentives. Indeed, by decreasing ex post payoffs, liability chills research ex ante. Thus, liability tends to be preferred for low externalities, while ex post regulation dominates for higher level of externalities. This result explains the difference between drugs (ex ante approval, high externality) and less risky products (liability, low externality) and also justifies why in practice liability and ex post regulation are rarely observed together.

The tradeoff between ex ante and ex post incentives is similar for all three instruments. When using a single instrument, the planner controls either the withdrawal standard (with liability or ex post regulation) or the adoption standard (with ex ante regulation). Moving away from the

[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.

socially optimal level of the controlled standard implies a second-order loss that is traded off against first-order gains in the other standards. If there are conflicts between the first-order effects, the optimal policy depends on the prior belief q_0, determining whether providing incentives for ex ante experimentation is the key dimension.

Consider the case of ex ante regulation in isolation. When the adoption standard is set at the socially optimal level S^*, the firm never withdraws, an extreme version of insufficient withdrawal incentives. The level of the externality determines whether incentives to experiment are excessive or insufficient. If e is low, since the firm never withdraws and does not suffer in the bad state, experimentation is excessive. In this case, setting a tougher approval standard than in the first-best solution is unambiguously socially beneficial. For sufficiently high e, however, experimentation incentives become insufficient. Being tougher when approving is beneficial ex post because once approved, a product is never withdrawn, but it is also costly ex ante because it decreases further experimentation incentives. The tradeoff between these two effects is resolved as a function of the prior belief.

Section 5.3 turns to our central question of the optimal combination of instruments. Our main result is that whenever ex ante or ex post regulation are used, they are always more lenient than the planner benchmarks. This is consistent with the regularities identified in the case of drugs and patents. The idea is the following. If the externality is low, the planner achieves the first best by using a combination of all three instruments. In this case the liability is used in equilibrium to limit excessive research incentives. On the contrary, when the externality is large, the first best is no longer achieved; liability rates are set at zero and the planner, in order to encourage research, commits to being more lenient than the socially optimal levels both in approval and withdrawal, regardless of the initial belief q_0.

Our model does not allow the planner to subsidize research. This assumption is motivated by two observations. First, the agencies in charge of approval and withdrawal regulation (FDA for drugs, USPTO for patents) are typically not responsible for choosing subsidies.[5] Second, in several applications there is in fact no subsidy program in place. In the absence of subsidies, for very large externalities, the main concern is to encourage experimentation; liability then is not used, as it chills research, and both ex ante and ex post regulation are weak. Moreover our results show that there is a critical level of externality at which the planner switches from using liability and socially optimal ex ante and ex post standards, to setting liability to zero in combination with weak regulation both ex post and ex ante. This critical level of externality is such that when the

[5] Viscusi, Vernon, and Harrington (1995, pages 785-786) highlight that an issue is the lack of coordination between regulatory and liability efforts. As they explain "these are two different institutional mechanisms that affect similar classes of economic concerns. In some cases, the companies are hit twice by these institutions." The issue is also relevant regarding subsidy programs.

planner chooses the approval and withdrawal standards at the socially optimal levels, the firm responds by abandoning research as the planner would do.

Section 6 considers the case in which the firm collects private (rather than public) information in the ex ante experimentation phase. We assume that the firm can make any report but is fined with a probability increasing in the size of the lie if the state turns out to be bad. As we show, if liability cannot be used, as for generic drugs or medical devices where ex ante approval shields from litigation, the planner is better off when evidence is privately observed by the firm than when it is public. Indeed, the fine for misreporting will be optimally set by the planner in such a way as to tolerate lying in equilibrium. The firm expects to pay with some probability a penalty in the bad state, thus chilling research incentives, an effect that is socially beneficial if the externality is small. The fine for misreporting serves as a substitute for liability.

Related Literature. Our contribution straddles the theoretical literature on optimal experimentation and the more applied law and economics literature on safety regulation and liability.

Following up on Wald's (1945) seminal work, the literature on sequential hypothesis testing focused on information acquisition before taking an irreversible action. Our ex-ante experimentation phase gains analytical traction by building on the continuous-time version of Wald's model developed by Dvoretsky, Keifer, and Wolfowitz (1953), Mikhalevich (1958), and Shiryaev (1967); see also Gul and Pesendorfer (2012), Chan, Lizzeri, Suen, and Yariv (2018), Henry and Ottaviani (forthcoming), and McClellan (2017) for economic applications building on the same workhorse.[6] Carpenter (2004) pioneered the application of experimentation models to approval regulation, focusing on the case with irreversible approval; see also Carpenter and Ting's (2007) analysis of how the firm can signal (private information about) quality to the regulator through the submission time.[7]

To an ex ante phase of experimentation à la Wald before adoption, we add a phase of ex post learning, which can eventually lead to withdrawal, effectively reversing the adoption decision. Learning in the ex post phase takes place only if the safe arm (represented by the withdrawal

[6]For applications to product search see also Branco, Sun, and Villas-Boas (2010) and Ke, Shen, and Villas-Boas (2016), and Ke and Villas-Boas (2018). Moscarini and Smith (2001) extend Wald's model to allow the decision maker to undertake multiple experiments in each instant. We instead allow for different experimentation phases, but still with a single experiment in each instant. Within a setting with a single irreversible decision, Che and Mierendorff (2017) push Wald's framework by allowing the decision maker to choose signal structures that favor learning about a state over the other.

[7]Carpenter and Ting (2007) in a discrete-time setting, as well as Carpenter (2004) in a continuous-time setting, model regulatory approval as a single-phase experimentation problem, without possibility of withdrawal. Carpenter (2004) justifies this assumption the fact that the professional regulator is concerned about the accuracy of the initial approval decision made. He shows that larger and older firms might receive a more favorable regulation, in terms of quicker approval, even when the regulator is not captured by producers.

decision) is not pulled, as in the bandit literature.[8] We isolate this first contribution—blending Wald with bandit—in our (decision-theoretic) social benchmark analyzed in Section 2. Given our interest in regulatory issues, the analysis in the body of the paper focuses on the case in which adoption can be reversed once at no cost. Appendix B analyzes in a self-contained way the more general case in which reversion of adoption is costly, thus recovering the classic Wald specification when reversion becomes prohibitively costly.[9]

Beyond characterizing the optimal balance between ex ante experimentation and ex post monitoring in the decision-theoretic benchmark, the bulk of the paper applies the framework to the strategic problem of dynamic regulation of an activity with uncertain externality. This second, more applied, contribution appears to be novel in the law and economics literature, in spite of its many applications.[10]

Related work in law and economics mostly focused on understanding the optimal way of incentivizing firms to take precautions so as to limit the negative externalities generated by risky activities. 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. Fiehe 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.[11] Our model focuses on settings with an externality of uncertain sign; thus we deal with the concern that activities with positive externalities might be chilled. Closer to our setting, Ottaviani and Wickelgren (2009) and (2011) offer a complementary modeling approach in the context of a two-period model exploring also signaling of the firm's private information, mostly in the context of competition policy and merger control.[12] Schwartzstein, and Shleifer (2013) show that when social returns to activity are higher

[8] See Bergemann and Välimäki (2008) for classic references and an introductory treatment of bandit models. 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 insightful 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).

[9] The literature on real options and investment under uncertainty—spearheaded by Arrow and Fisher (1974) and Henry (1974) and overviewed by Dixit and Pindyck (1994)—focuses on the impact of exogenous information flow on the incentives for adoption. Instead, our model endogenizes the flow of information either based on dry bench experimentation (with a state-independent cost) or wet bench learning (with an implicit cost determined by the state-dependent payoff associated with the adoption decision).

[10] Orlov, Skrzypacz, and Zryumov (2018) analyze dynamic information control by an agent who aims at persuading a principal to delay withdrawal—the simpler information structure we posit allow us to analyze information arrival *both* before and after adoption up to withdrawal.

[11] See Daughety and Reinganum (2013) for a broader overview of the economic analysis of products liability.

[12] In the area of competition policy Rey (2003, Section 4.2) also informally discusses the pros and cons of ex ante regulation v. ex post antitrust.

than private returns, liability chills economic activity opening room for regulation.

In an informal discussion, Viscusi, Vernon, and Harrington (1995, pages 785-786) raise precisely the issues our paper studies theoretically. They describe 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 in the context of our stylized model.

Our model abstracts away from contracting between the firm and consumers, which is instead the focus of the literature on product recalls; see, for example, Welling (1991), Marino (1997), and Rupp and Taylor (2002). In particular, Spier (2011) shows that, even under strict liability—whereby the firm needs to fully compensate any harm caused—the buyback price is inefficiently low, because of the firm's monopsonistic position in the ex post stage when products are recalled.

In a recent contribution, Hua and Spier (2018) show that firms with market power find it optimal to underprovide safety and to disclaim responsibility for consumer harm, provided that consumers who use the product more intensely are also more likely to suffer harm. This result offers a rationale for liability regulation, even when firms set prices and are allowed to reimburse the damages created by the products they sell. Instead, our work does not consider the interaction between the firm and consumers. Taking as a given the need for regulation, we compare different forms of intervention and characterize the optimal regulatory mix of instruments.

2 Model

Model. Two risk neutral players, a firm i and a social planner p, interact in continuous time under uncertainty about the state of the world θ, which can be either good G or bad B.[13] The game is divided in two phases. In the first phase, at each instant of time a decision needs to be made, either reject R, experiment E, or adopt A. Rejection is irreversible: the game ends following R and players obtain 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 active player can either continue C (staying on the market) or withdraw W (ending the game).

Payoffs are collected only during the ex post phase, except for the cost of experimentation. While on the market (i.e., following adoption and up until withdrawal), the firm collects a flow payoff v_i, independent of the state θ.[14] The social planner, instead, collects flow state-dependent

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

[14] A justification for this assumption is that consumers don't have access to the information about the state post approval and thus the price of the product does not change as a function of the state.

payoffs $v_p^G \geq v_i > 0$ in the good state and $v_p^B < 0$ in the bad state. Thus, the firm's activity generates a positive externality $v_p^G - v_i \geq 0$ in the good state and a negative externality $v_p^B - v_i < 0$ in the bad state on the rest of society. For the purpose of comparative statics, it is useful parametrize the firm's (private) payoff as $v_i = (1-e)v_p^G$, where the externality rate $e \in [0,1]$ can be interpreted as the fraction of the planner (social) payoff that the firm does not capture when the state is good.

All players initially share the same prior about the state $q_0 = \Pr\{\theta = G\}$ and discount future payoffs at the same rate $r \geq 0$.

Information Arrival. We now describe information arrival in the two phases, $n \in \{1,2\}$, where $n = 1$ denote the ex ante experimentation phase and $n = 2$ the ex post learning phase.

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. Acquiring information over a period of time dt in the ex ante phase costs the firm cdt while it is costless in the ex post phase.

The realization of the stochastic process at time $t > 0$ is a sufficient statistic for all the information collected until this instant of time and will be used to update beliefs. We express beliefs in terms of the log-odds ratio

$$\sigma_t = \log \frac{q_t}{1-q_t} \tag{1}$$

and denote $\mathsf{S} = \log \frac{S}{1-S}$, $\mathsf{s} = \log \frac{s}{1-s}$, and $\mathsf{z} = \log \frac{z}{1-z}$ respectively the log-odds of the adoption, rejection, and withdrawal standards, soon to be introduced.

3 Planner Benchmark: Balancing Ex Ante Experimentation and Ex Post Learning

As a benchmark for the rest of our analysis, we consider the decision of the social planner when in charge of all decisions in both phases:

Proposition 1 *The solution of the social planner consists of three standards, s^* (the rejection standard), S^* (the adoption standard), and z^* (the withdrawal standard), such that:*

(a) *In the ex ante phase, the planner:*

- **(i)** *stops experimentation and rejects if $q \leq s^*$,*
- **(ii)** *experiments if $s^* < q < S^*$, and*
- **(iii)** *stops experimentation and adopts moving to the ex post phase if $q \geq S^*$;*

(b) *In the ex post phase, the planner withdraws if $q \leq z^*$;*

(c) *Furthermore all standards are independent of the current belief q and are such that $s^* \le z^* \le S^*$.*[15]

Proposition 1 characterizes the optimal balance between ex ante experimentation and ex post learning. Even though information is revealed at no direct cost in the ex post phase, an indirect cost is nevertheless incurred when the expected flow payoff is negative (because the belief favors the bad over the good state) so that it would be myopically optimal to withdraw. Next, we describe the mechanics of the model, starting from the last of the two phases.

3.1 Ex Post Learning: Bandit Problem

Consider first the ex-post phase. Assume that the belief has reached the adoption threshold S, so that the ex ante phase resulted in adoption. The ex post problem can be seen as a bandit in which at each instant of time the planner chooses between a safe arm (withdrawal W with payoff zero) or a risky arm (continue C with a state dependent flow of profit whose value is uncertain). This problem is equivalent to a one-time one-option optimal stopping problem.[16]

The expected payoff in the second phase can be expressed as

$$\hat{u}_p^2(\mathsf{S}) = \frac{e^{\mathsf{S}}}{1+e^{\mathsf{S}}}\left(\frac{v_p^G}{r}\right)(1-\psi(\mathsf{S},G,\mathsf{z}^*)) + \frac{1}{1+e^{\mathsf{S}}}\left(\frac{v_p^B}{r}\right)(1-\psi(\mathsf{S},B,\mathsf{z}^*)),$$

where $\psi(\sigma,\theta,\mathsf{z}^*)$ is the expected discounted probability of reaching withdrawal threshold z^*in state θ starting at σ; closed-form expressions for $\psi(\sigma,\theta,\mathsf{z}^*)$ are presented in the proof of Proposition 1 (following Stokey 2009).

The dashed-dotted blue line in Figure 1 shows the ex post value function and the optimal solution z^* of the bandit problem. The value is positive for $q \ge z^*$, zero for $q \le z^*$, and tangent to the zero horizontal line exactly at z^*, i.e. satisfies the smooth-pasting condition. The dotted line in Figure 1 corresponds to the expected value of immediate approval without withdrawal.

3.2 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 denoted $\hat{u}_p^2(\sigma_\tau)$ if the belief at adoption is σ_τ. In the ex ante phase, this becomes a standard Wald problem where the planner abandons for sufficiently bad news ($q \le s^*$), experiments for intermediate beliefs ($s^* < q < S^*$), and adopts following good news ($q \ge S^*$). Differently from the classic Wald problem, the payoff upon adoption (i.e. the ex

[15]Notice that the withdrawal standard z^* will not depend on the belief S^* at which adoption took place.

[16]It can be shown that whenever it is optimal to pull the safe arm then it is optimal to keep pulling it afterwards. The optimal withdrawal must be an irreversible decision.

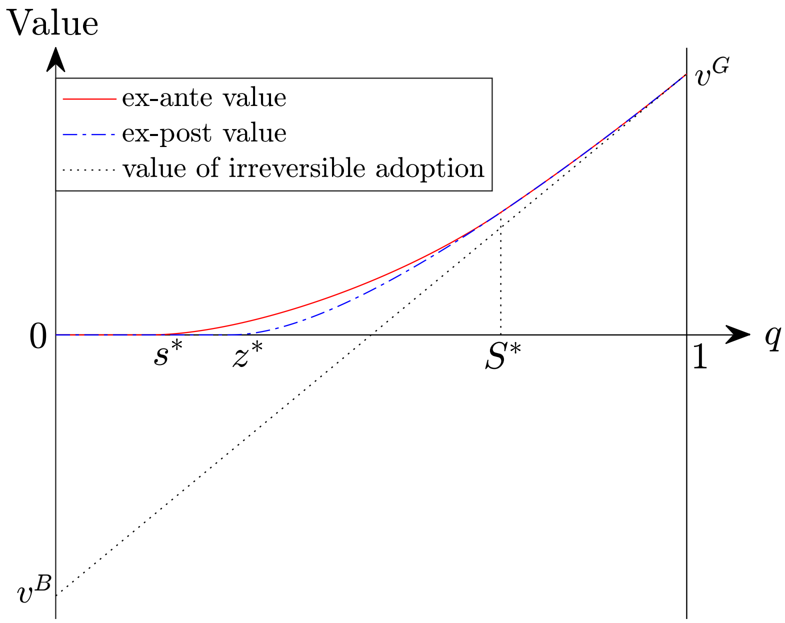

Figure 1: Value functions in the different phases.

post value) is now endogenous. Adoption at a higher S yields a higher expected payoff from the ex post phase for two reasons: first, the expected flow benefit is higher at the start and, second, eventual withdrawal becomes less likely. The expected payoff can be written as

$$\begin{aligned}\hat{u}_p^1(\sigma_0) \quad = \quad & -\frac{c}{r}+\frac{e^{\sigma_0}}{1+e^{\sigma_0}}\left[\left(\hat{u}_p^2(\mathsf{S}^*,G)+\frac{c}{r}\right)\Psi(\sigma_0,G,\mathsf{s}^*,\mathsf{S}^*)+\frac{c}{r}\psi(\sigma_0,G,\mathsf{s}^*,\mathsf{S}^*)\right] \\ & +\frac{1}{1+e^{\sigma_0}}\left[\left(\hat{u}_p^2(\mathsf{S}^*,B)+\frac{c}{r}\right)\Psi(\sigma_0,B,\mathsf{s}^*,\mathsf{S}^*)+\frac{c}{r}\psi(\sigma_0,B,\mathsf{s}^*,\mathsf{S}^*)\right],\end{aligned}$$

where $\Psi(\sigma,\theta,\mathsf{s},\mathsf{S})$ and $\psi(\sigma,\theta,\mathsf{s},\mathsf{S})$ are the expected discounted probabilities of reaching respectively the adoption S and rejection s standards in state θ starting at belief σ (following Stokey 2009).

The withdrawal standard used in the ex post phase plays an important role in the ex ante phase. Indeed, z^* is the belief $q=z^*$ at which in the ex ante phase the planner is indifferent between rejection (yielding 0 payoff) and adoption (as going to the ex post phase leads to immediate withdrawal yielding 0). Therefore, in the ex ante phase, around z^* experimentation has an option value—thus, the region of ex ante experimentation includes the withdrawal standard: $s^* \leq z^* \leq S^*$.[17]

Given z^*, Figure 2 plots the dashed-dotted lower best reply $b_p(S)$ (optimal choice of s for a given adoption standard S) and the continuous upper best reply $B_p(s)$ (optimal choice of S for a given rejection standard s) in the regular belief space.[18] These two best replies are respectively

[17]Given z^*, the pair $(\mathsf{s}^*,\mathsf{S}^*)$ solves $\hat{u}_p^1(\sigma_0)\equiv\max_{\mathsf{s},\mathsf{S}} u_p^1(\sigma_0,\mathsf{s},\mathsf{S})$.

[18]See Appendix B for the construction of $b_p(S)$ and $B_p(s)$.

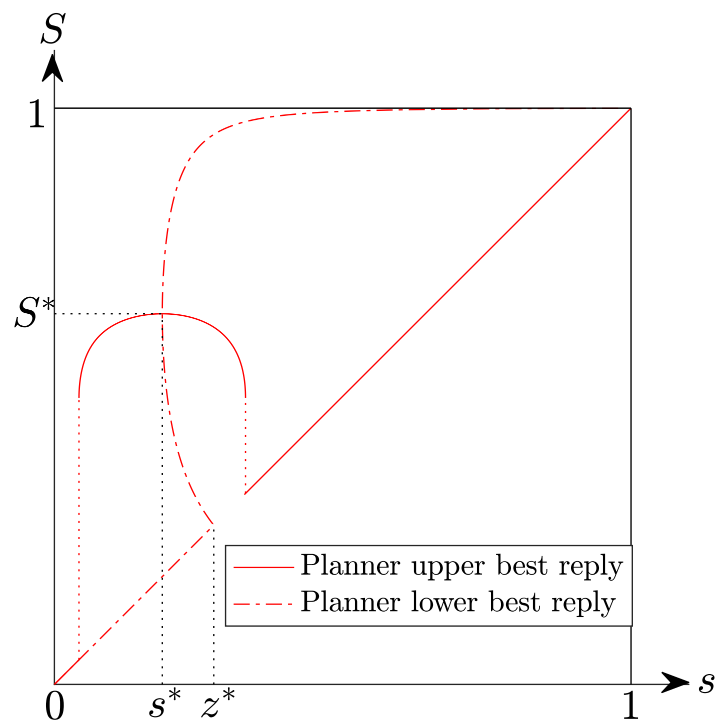

Figure 2: Social planner's best replies.

defined implicitly by the two first-order conditions (2) $\partial u_p^1/\partial s = 0$ and (3) $\partial u_p^1/\partial S = 0$ presented in Appendix A. The optimal ex ante standards (s^*, S^*) are then defined at the intersection of the two best replies.

As can be seen in Figure 1, the possibility of conducting research in the ex ante phase increases the expected welfare of the planner. In particular, around z^* there is value of information while immediate approval at z^* gives a zero payoff. The ex ante value is tangent to the ex post value (dashed-dotted blue) exactly at s^* and S^*, i.e. the pair (s^*, S^*) solving the smooth-pasting conditions. Overall, the planner strikes a balance between ex ante experimentation and ex post monitoring.

4 Regulatory Environments

In practice, three main instruments are typically employed to align private incentives with the social planner's objectives: liability, approval (ex ante) regulation and withdrawal (ex post) regulation. Before comparing the performance of these instruments in the context of our model in Section 5, this section outlines the institutions in a number regulatory environments, focusing mostly on our leading applications to drug safety regulation and to patent grants, in particular in the case of the US.

Drug and Medical Devices. For pharmaceutical products, the key state variable that determines market introduction is their effectiveness relative to the risk of serious side effects. The firm and the regulator can collect evidence through clinical trials in the ex ante phase and through surveillance in the ex post phase.[19]

- *Ex ante regulation*: To sell a new prescription drug, the sponsor must provide a series of codified clinical trial results based on which the FDA can allow market introduction. The FDA and the sponsor must also agree on the labelling of the product.

- *Ex post regulation*: Post approval, pharmaceutical firms are required by the FDA to report adverse events experienced by patients or results of additional clinical trials they conduct. Patients and doctors can also directly report to the FDA through the MedWatch system. The FDA can at any point request a change in the labeling of the drug or can order the firm to recall the drug.[20]

- *Liability*: Product liability law in the case of drugs, like in the case of most products, imposes legal obligations to compensate people injured by the product. The defect most commonly invoked for pharmaceuticals is the warning failure, whereby the manufacturers is liable for failing to warn about risks they knew or should have known about.[21] In most cases that led to large levels of liability (such as the Vioxx case), the defining feature is that the companies failed to report to the FDA or to the medical community evidence they accumulated post marketing.[22] A long lasting debate, as mentioned in the introduction is whether pharmaceutical companies should be protected from liability if they complied with FDA regulations—a defense that is valid for generic drugs and medical devices. The Supreme Court ruled against this defense in *Wyeth v. Levine*, even though certain states adopt this approach.[23]

[19]See Daemmrich (2007) for a historical account of the evolution of pharmacovigilance in the US.

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

[21]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.

[22]It was judged in the Vioxx case that the warnings to physicians were inadequate in light of what Merck knew at the relevant times about cardiovascular risks from their own studies as well as published or unpublished literature.

[23]In PLIVA, *Inc. v. Mensing*, 131 S. Ct. 2567 (2011), the US Supreme Court rules that generic drug manufacturers are not liable for injuries arising out of their failure to warn of dangers from the use of their drug, provided that they use FDA required warnings. Thus, generic drug manufacturers are sheltered from the main threat of liability. See Friedman and Wickelgren (2017) for an economic analysis.

Lawsuits against pharmaceutical companies are much more common in the US than in other countries around the world. For instance, Lybecker and Watkins (2015) report that "there is no evidence of any court cases with positive settlement payments documented within the UK." The essential difference is that a UK drug manufacturer can avoid litigation by arguing of having been unaware of the side effect.[24] At the same time, prices of drugs are much higher in the US. We discuss in the conclusion how our model can explain this combination of features that characterizes the US.

In conclusion, the main regulatory tool for drugs remains ex ante regulation with strict clinical trial testing. Product liability, in particular due to failure to warn, still plays a role at least in the US, but in spite of the Supreme Court ruling remains partly preempted by ex post regulation. Ex post regulation does exist but tends to be still rather weak.

Patent validity. In the case of granting patents the state of the world is whether the patent is valid or not: a valid patent provides a positive payoff to society, while an invalid patent causes social harm as the patent holder can sue productive firms that are supposedly infringing.[25] As Hall et al. (2004) frame it: "Low-quality patents can create considerable uncertainty among inventors or would-be commercializers of inventions and slow either the pace of innovation or investment in the commercialization of new technologies."

- *Ex ante regulation*: The US patent office (USPTO) determines whether the patent meets the three criteria—novelty, usefulness and non obviousness—necessary for being approved. Many scholars (Jaffe and Lerner 2007) believe that in the US, this ex ante regulation is rather weak, in the sense that many invalid patents are being granted, partly because the patent office is overcrowded, partly also because the incentives of the patent office used to be biased in favor of approval.[26]

- *Ex post regulation*: If the patent holder sues for infringement, the defendant can adopt an invalidity defense and attempt to show that the patent is invalid. If the defendant wins that case, the patent is definitively withdrawn. This ex post regulation is considered weak for two reasons. First, as reported by Ford (2013), the defendant has more incentives to use a noninfringement defense rather than an invalidity defense as it cannot benefit from the positive externality it provides to others by having a bad patent invalidated. Second,

[24] 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."

[25] Galasso and Schankerman (2015) provide causal evidence that patents can decrease cumulative innovation.

[26] Ford (2013) reports that recently the USPTO has received more than five hundred thousand applications per year and granted almost half as many.

there is a strong burden of proof on the defendant to prove a patent invalid; according to Hall et al. (2004), "patents are born valid, thus enjoying a presumption of validity during the court proceedings. Furthermore, the evidentiary standard for proving that a claim is invalid is clear and convincing evidence, a standard considerably higher than the mere preponderance of proof required in the typical civil suit."[27] Since 2012, a plaintiff may also challenge one or more claims of a patent directly in front of the Patent Trial and Appeal Board, a body of the USPTO (the procedure is called Inter Partes Review). The procedure also puts the burden of proof on the petitioner.

- *Liability*: No liability can be imposed on the holder of a bad patent that is proven invalid. Even though courts intervene ex post to potentially rule patents as invalid, firms that paid royalty fees to the patent holders prior to invalidation cannot claim financial compensation.

In comparison with the US, the European Patent Office appears to be stricter in granting patents (ex ante regulation). A similar system of ex post invalidation by courts is in place in Europe. It is hard to determine whether overall the ex post regulation is stricter in Europe or in the US, but there is one dimension that limits its scope in Europe: lawsuits for infringement and thus counter lawsuits for invalidity are brought in front of national courts, so that a patent can be ruled invalid in one country while still being valid in another.[28]

In conclusion, in the case of patent awards, there is a mix of ex ante and ex post regulation, which tend to be weak, especially in the US. Liability is not part of the mix used by the planner. In the above discussion, the issue at stake is the validity of the initial patent.

Other Applications. The mix of regulatory instruments employed varies greatly across settings beyond drugs and patents. For safety regulation of products other than pharmaceuticals for which the risk of negative externalities is more limited, the main instrument is liability and regulation plays a minor role. In the case of projects with environmental risk (oil drilling in sensitive areas or introduction of genetically modified organisms) or merger policy, withdrawal is very costly; ex ante regulation is the natural instrument.[29] Finally, road safety is regulated both ex ante (requiring a driver's license) and ex post (withdrawing the license following a series of traffic offenses) as well as through liability following accidents.[30]

[27]The law requires that patents be presumed valid and that a party asserting that a parent claim is invalid bears the burden of proof.

[28]Coyle (2012) describes a case where a patent was invalidated in United Kingdom, France, Belgium, and Austria, yet upheld in Germany and the Netherlands.

[29]Appendix B extends our baseline model to introduce reversibility costs.

[30]For applications in the area of consumer financial protection see Posner and Weyl (2013).

5 Approval Regulation, Withdrawal, and Liability

We now analyze the performance of the different regulatory instruments that can be used to align the firm's incentives and confront the model's predictions with a number of regularities that emerge from the description in Section 4. We initially focus on situations in which at date zero the planner controls only one of the three tools in isolation. We model the three tools as follows:[31]

- **Liability.** The planner commits at date zero to a liability rate $L \in \left[0, v_i - v_p^B\right]$ that is imposed on the firm per unit of time on the market if $\theta = B$.[32] In state B, the firm thus obtains payoff $v_i - L$; in state G, instead, the payoff remains unaffected at $v_i = (1-e)v_p^G > 0$. Under the liability regime, the firm controls rejection $s_{ii}(L)$, adoption $S_{ii}(L)$, and withdrawal $z_{ii}(L)$.
- **Ex Ante Regulation.** The planner commits at date zero to an approval standard S_{pi}^*. The firm controls rejection s_{pi} and withdrawal z_{pi}.
- **Ex Post Regulation.** The planner commits at date zero to a withdrawal standard z_{ip}^* in the ex post phase. The firm fully controls the ex ante phase by setting rejection s_{ip} and adoption z_{pi}.

Studying the instruments in isolation allows us to introduce and compare the key forces at play. Section 5.3 then characterizes the optimal combination of the three instrument: the planner commits at date zero to a combination $(S_{pp}^*, z_{pp}^*, L_{pp}^*)$, while the firm only controls rejection s_{pp}.

5.1 Liability v. Ex Post Regulation

The first regularity that emerges from the discussion in Section 4 is that liability (used for instance for regular product introductions) and ex post regulation (used for drugs and patents) are rarely used in conjunction. We therefore start by comparing these two instruments.

Liability. Under a strong liability regime, defined as a regime where the liability rate is set at $L = \overline{L} = v_i - v_p^B$ in a way to fully compensate the damage caused, the incentives of the firm and the social planner are by construction perfectly aligned in the bad state. In the absence of externalities, $e = 0$, strong liability naturally achieves the first best since the firm and the social

[31] Given a threshold $x \in \{s, S, z\}$, we denote by x_{ab} the threshold resulting in the regulatory environment where player a controls ex ante adoption and player b controls ex post withdrawal.

[32] We purposefully abstract away from the details of the judicial procedure by assuming that courts impose an expected penalty equal to L per unit of time in the market conditional on state B.

planner obtain the same payoff in all states. If $e > 0$, strong liability no longer induces the first best: since the firm's payoff in the good state is socially insufficient, the firm rejects too early, $s^* < s_{ii}\left(\overline{L}\right)$, adopts too late, $S^* < S_{ii}\left(\overline{L}\right)$, and withdraws too early, $z^* < z_{ii}\left(\overline{L}\right)$. The social planner should constrain the courts to be more lenient.

The planner can in fact choose the liability rate L to induce any withdrawal standard. In particular, we denote by $\hat{L}$ the liability rate that induces the ex post optimal withdrawal, $z_{ii}(\hat{L}) = z^*$. Under this liability rate, as illustrated in Figure 3, the firm's lower best reply $b_i(S)$ (dashed-dotted blue) is to the right of the social planner's $b_p(S)$ (dashed-dotted red), since the firm expects a lower payoff from adoption for a given set of realized standards. The firm's upper best reply $B_i(s)$ (continuous blue) is below the social planner's $B_p(s)$ (continuous red), given that the firm attaches a lower value to information. Overall, the firm rejects too early, $s^* < s_{ii}(\hat{L})$, and adopts too early $S_{ii}(\hat{L}) < S^*$.

The planner can achieve a higher level of social welfare by committing to a liability rate different from $\hat{L}$. From the ex post perspective, this increase in liability induces a second-order loss, as withdrawal is moved away from the optimal ex post standard z^*. Deviating from that level, however, may generate a first-order gain in the change of ex ante actions that the firm is induced to take. Recall that the planner wants to encourage research at the bottom (decrease s_{ii}) and delay approval (increase S_{ii}). Given that any change that decreases s_{ii} also decreases S_{ii}, the planner must trade off conflicting forces in terms of ex ante actions.

Varying z around z^*, has conflicting effects on the rejection and adoption standards. By choosing a liability rate strictly above $\hat{L}$, the value of information for the firm is decreased, leading the firm to adopt later and reject earlier. If the rate is chosen strictly below $\hat{L}$, the expected payoff upon approval is increased, leading the firm to adopt earlier and reject later. Overall, the planner optimally commits to a liability rate below $\hat{L}$ if the starting belief q_0 is low and the main concern is to encourage experimentation. On the contrary if q_0 is high, the main concern of the planner is to discourage early adoption and the planner should commit to a liability rate strictly above $\hat{L}$.

Proposition 2 *If $e > 0$, optimal liability L^* is interior $0 < L^* < \overline{L}$ and such that:*

(a) *The firm rejects too early $s_{ii}(L^*) > s^*$ and adopts too early $S_{ii}(L^*) < S^*$;*

(b) *L^* is decreasing in e and function of initial belief q_0: $\exists\, \hat{q} \in (\underline{q}, \overline{q})$ such that:*

- **(i)** *for any $q_0 \in (\underline{q}, \hat{q})$, the planner commits to low liability $L^*(q_0) < \hat{L}$,*
- **(ii)** *for any $q_0 \in (\hat{q}, \overline{q})$, the planner commits to high liability $L^*(q_0) > \hat{L}$.*

Proposition 2 implies that, in a system that relies exclusively on liability such as regulation of product safety, the optimal liability rate is a function of the level of externality, a question

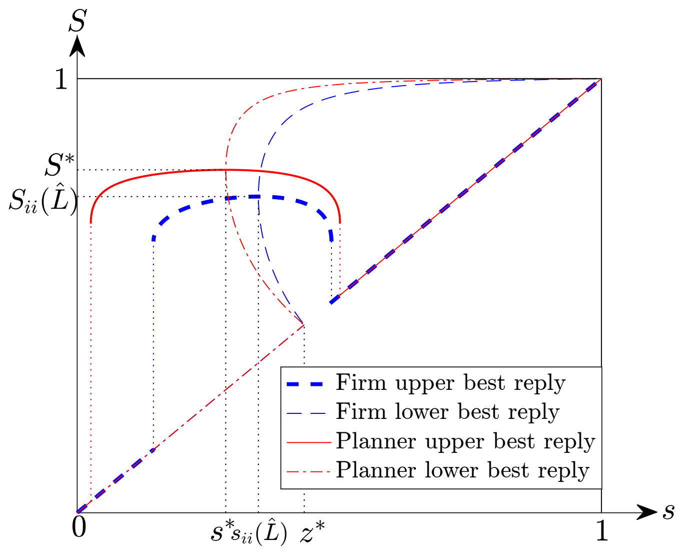

Figure 3: Best replies of social planner and firm under liability.

we return to in the next section, but also of the initial state of knowledge at the start of the experimentation process (represented by q_0), that varies in practice for instance depending on how novel the product is.

Ex Post Regulation. Before comparing ex post regulation with liability, we derive the properties of optimal regulation, when the planner commits to a withdrawal standard z^*_{ip} at date zero. Optimal regulation, as in the case of optimal liability characterized in Proposition 2, rests on the logic that it is socially optimal to incur a second-order loss by moving away from the optimal standard under control (withdrawal standard) against a first-order gain in the other standards.

Proposition 3 *The optimal withdrawal standard z^*_{ip} is a function of q_0, with $\bar{e}_{ip}$ such that:*

(a) *If $e \leq \bar{e}_{ip}$ the planner commits to be* ***tougher*** *than the socially optimal withdrawal level: $z^*_{ip} > z^*$;*

(b) *If $e > \bar{e}_{ip}$, there exists $\bar{q}(e)$ such that:*

- **(i)** *if $q_0 < \bar{q}(e)$, the planner commits to be* ***more lenient*** *than the socially optimal withdrawal level: $z^*_{ip} < z^*$,*
- **(ii)** *if $q_0 \geq \bar{q}(e)$, the planner commits to be* ***tougher*** *than the socially optimal withdrawal level: $z^*_{ip} > z^*$.*

As in the case of liability, the planner might want to commit to a withdrawal standard different from the ex post optimal standard z^* in order to influence ex ante incentives. When e is small, the firm, since it does not incur a negative payoff in the bad state and has a payoff in the good

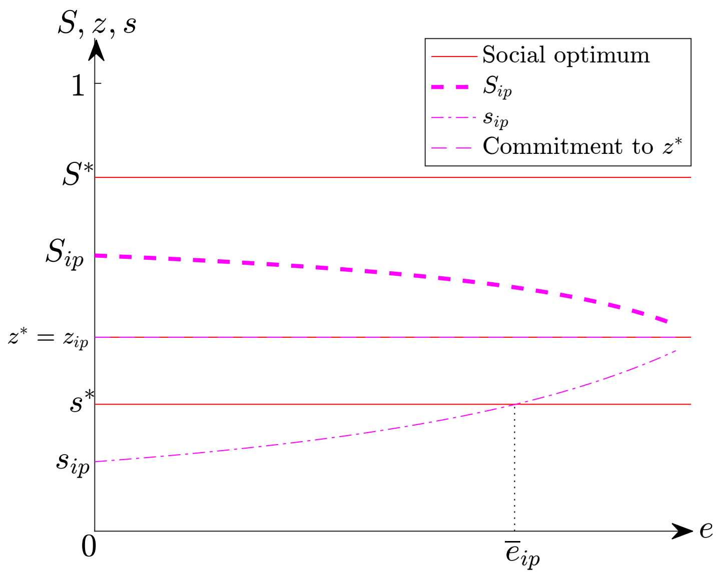

Figure 4: Commitment to $z_{ip} = z^*$ and induced ex ante incentives as a function of e.

state similar to that of the planner, has excessive incentives both to approve ($S_{ip}(z^*) < S^*$) and to experiment ($s_{ip}(z^*) < s^*$). Thus, when e is small, it is unambiguously optimal to increase z_{ip} above z^* to chill experimentation incentives, thus pushing the firm to reject earlier and approve later.

When e is larger, it is no longer necessary to chill research at the bottom. We plot in Figure 4 as a function of e, both the socially optimal standards s^*, z^*, S^* and the induced ex ante standards s_{ip} and S_{ip} if the firm commits to a withdrawal level $z_{ip} = z^*$. Whereas socially optimal standards are independent of e, in the decentralized problem, as e increases, the firms both abandons and adopts more quickly. In particular, there exists a benchmark value $\overline{e}_{ip}$ such the firm chooses $s_{ip} = s^*$: the externality is sufficiently large that, even though the firm does not obtain a lower payoff in the bad state, the low payoff in the good state implies that the incentives to keep experimenting are sufficiently chilled. If the externality is above this benchmark value, $e > \overline{e}_{ip}$, it becomes necessary for the planner to encourage experimentation. The withdrawal standard z_{ip} then has two conflicting effects on the ex ante standards s_{ip} and S_{ip}. Lowering z_{ip} below z^* entails the benefit of encouraging experimentation by delaying rejections, but also the drawback of speeding up adoption. The optimal balance between these two effects depends on the starting belief q_0. If the prior is sufficiently low, the most pressing concern is to encourage experimentation, thus pushing the planner to be more lenient. If instead q_0 is high, adoption incentives are more important, thus inducing the planner to be tougher.

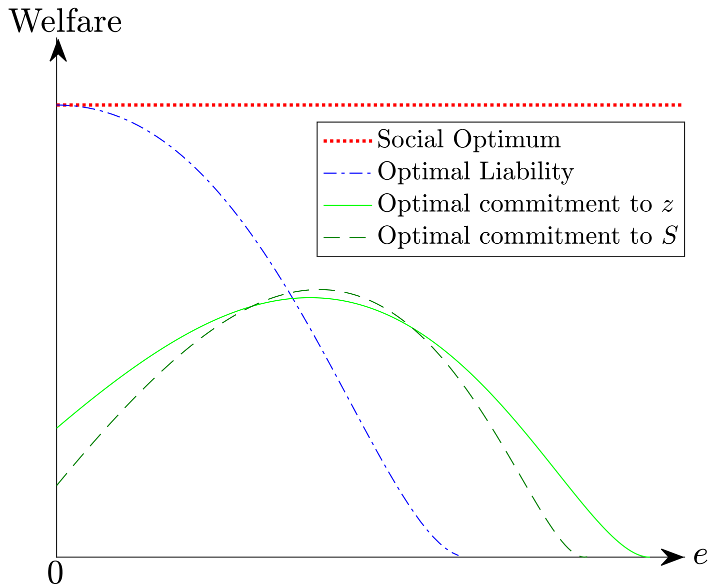

Figure 5: Welfare under optimal ex post commitment, optimal ex ante commitment, and optimal liability, taken individually.

Comparing liability and Ex Post Regulation. Liability can be set to induce the same withdrawal as ex post regulation: for any z^*_{ip}, there exists $L \in (0, \overline{L})$ such that the resulting withdrawal standard $z_{ii}(L) = z^*_{ip}$. For instance, in the case of drugs, 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. However, liability results in more powerful ex ante incentives than ex post regulation because it decreases payoffs and thus has a stronger chilling effect on experimentation, for the same ex post standard. Liability thus dominates for low e while ex post regulation dominates for high e.

Proposition 4 *There exists $\tilde{e}$, such that an optimal liability system is preferred to an optimal ex post regulation if and only if $e \leq \tilde{e}$.*

Figure 5 illustrates the welfare performance of all three instruments taken individually as a function of e, compared to welfare at the social optimum. In particular, liability yields higher welfare if and only if the externality is low enough (in particular achieves the social optimum for $e = 0$) and this is consistent with regularities observed in practice. When the externality rate is high, liability chills experimentation incentives too much. Ex post regulation, instead, has a softer effect on those incentives and hence welfare dominates liability, explaining why, this instrument is rarely used for the introduction of patents or drugs. On the contrary, when the externality rate is low, as in the case of safety regulation of regular products, liability is preferred.

5.2 Regulation: Lenient or Tough?

The second regularity that emerges from the discussion in Section 4 is that regulation, both ex ante or ex post, tends to be lenient. In the case of patents, ex post regulation takes the form of patent invalidation, which is relatively light touch, given that the burden of proof is put on the challenger. Similarly, in drug regulation there is limited monitoring of side effects for approved drugs through so-called post marketing studies.

As shown in Proposition 3, when the planner can only use ex post regulation, the commitment is weak only if e is large, i.e instances where experimentation needs to be encouraged, and q_0 is low, so that experimentation is a more pressing issue than approval. The exact same logic applies in the case of optimal ex ante regulation where the planner controls only the ex ante standard S^*_{pi}.

Proposition 5 *The optimal approval standard S^*_{pi} is a function of q_0, with $\bar{e}_{pi}$ such that:*

(a) *If $e \leq \bar{e}_{pi}$, the planner commits to be **tougher** than socially optimal adoption level: $S^*_{pi} > S^*$;*

(b) *If $e > \bar{e}_{pi}$, there exists $\bar{q}(e)$ such that:*

- **(i)** *if $q_0 < \bar{q}(e)$, the planner commits to be **more lenient** than the socially optimal adoption level: $S^*_{pi} < S^*$,*
- **(ii)** *if $q_0 \geq \bar{q}(e)$, the planner commits to be **tougher** than the socially optimal adoption level: $S^*_{pi} > S^*$.*

As in Proposition 5, starting from the socially optimal adoption standard S^*, the firm has insufficient withdrawal incentives and excessive (resp. insufficient) experimentation incentives if the externality level e is low (resp. high). Similar to the case described above, there is a tradeoff between a second-order loss from moving away from S^* versus a first-order gain in rejection standard s_{pi}.

The cutoff value $\bar{e}_{pi}$ is defined as the level of externality such that, if the planner commits to an adoption standard $S^*_{pi} = S^*$, the firm responds by choosing the socially optimal research standard $s^*_{pi} = s^*$. This is represented in Figure 6 where $\bar{e}_{pi}$ is the externality such that s_{pi} as a function of e crosses the horizontal line equal to s^*. By definition of $\bar{e}_{pi}$, if the externality is smaller than $\bar{e}$, increasing S^*_{pi} above S^* is beneficial both because it discourages excessive experimentation, increasing rejections and delays adoption, which is socially beneficial since the firm has no incentive to withdraw once the product is adopted. On the contrary, when $e > \bar{e}_{pi}$, experimentation incentives need to be encouraged. Changing S^*_{pi} has conflicting effects. If q_0 is small, the effect on experimentation incentives is most pressing and the planner encourages experimentation by reducing adoption standards $S^*_{pi} < S^*$. On the contrary, when q_0 is higher,

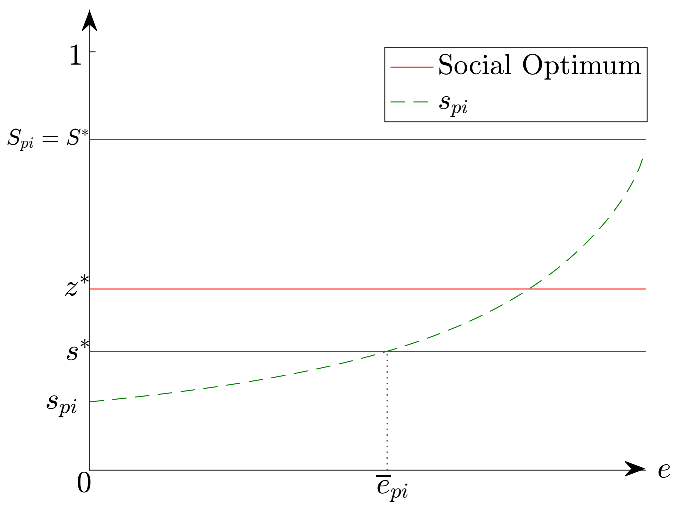

Figure 6: Research incentives of the firm when planner commits to $S_{pi} = S^*$.

the planner is more worried about the fact that the firm won't withdraw once approval has been granted and thus prefers to be tougher on adoption $S^*_{pi} > S^*$.

Overall, Propositions 3 and 5 give a nuanced view on the second regularity drawn from the evidence in Section 4. Whether the planner, who can control a single instrument, sets regulatory standards that are tougher or weaker than the socially optimal levels depends on the externality rate e and on the starting belief q_0.

5.3 Optimal Mix of Instruments

Now we show that when the planner controls all three instruments, regulation (ex post and ex ante) is always weaker than or equal to the socially optimal levels, in full accordance with the second regularity described above:

Proposition 6 **(a)** *The socially optimal mix of the three instruments $(S^*_{pp}, z^*_{pp}, L^*_{pp})$ is such that:*

- **(i)** *the first best is achieved if and only if $e < \hat{e}$, and all three instruments are used,*
- **(ii)** *liability is set to zero if $e \geq \hat{e}$.*

(b) *In an optimal mix of instruments, both* ***ex post and ex ante commitments are chosen equal or more lenient*** *than the socially optimal levels z^* and S^*.*

According to Proposition 6, there is a critical level of externality $\hat{e}$ such that, if e is less than $\hat{e}$, the planner chooses approval and withdrawal standards at the socially optimal levels and uses liability to chill research incentives. When instead e is greater than $\hat{e}$, the planner stops using liability and encourages research by being more lenient with both ex ante and ex post regulatory standards. Thus liability is only used when experimentation needs to be discouraged.

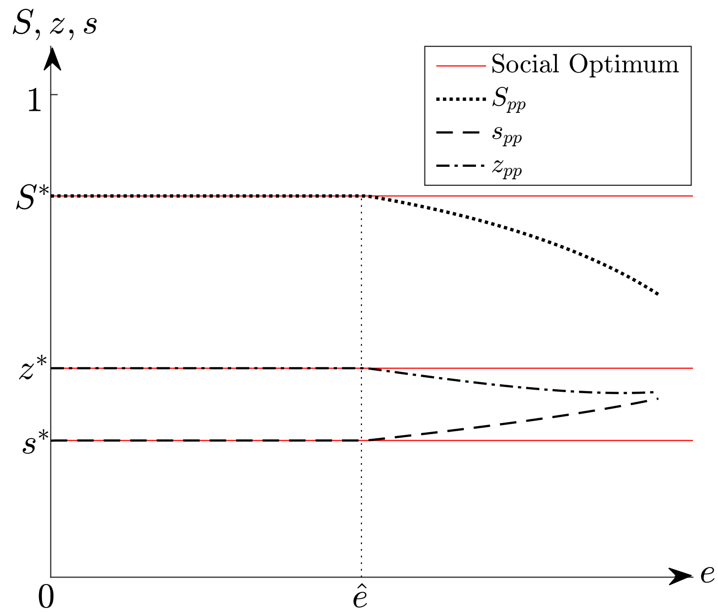

Figure 7: Optimal combination of instruments.

To understand how the critical level $\hat{e}$ is determined, consider the case where the planner commits to a combination of the socially optimal ex ante standard S^* and ex post standard z^*, but does not use liability. Given these commitments, the firm optimally decides on when to abandon experimentation, a decision we denote s_{pp}. For $e = 0$, incentives to experiment are excessive ($s_{pp} < s^*$) since the firm has the same payoffs in the bad and good state. For $e = 1$, on the contrary experimentation incentives of the firm are insufficient $s_{pp} > s^*$.[33] Since the lower best reply $b_i(S)$ increases when e increases, there exists a critical value $\hat{e}$, such that if $e = \hat{e}$ then $s_{pp} = s^*$. By definition, $\hat{e}$ induces the socially optimal rejection decision by the firm, so for that level of externality the first best can be achieved using only a combination of ex ante and ex post commitments.

If $e < \hat{e}$, the first best is no longer achieved by using solely a combination of the two regulatory instruments. The planner can use in addition liability that, for the same level of approval and withdrawal standards, discourages excessive experimentation. In the limit where $e = 0$, we already saw in Section 5.1 that strong liability alone achieves the first best. When $e = \hat{e}$ liability is set to zero and, as already highlighted, first best is achieved only with ex ante and ex post regulation. If, instead, $e \in (0, \hat{e})$ there exists an interior liability level that achieves the first best.

On the contrary, if $e > \hat{e}$, the first best is no longer achievable. In this case, when the adoption and withdrawal standards are chosen at the socially optimal levels, experimentation incentives are insufficient. Liability is no longer a useful instrument since it further discourages experimen-

[33] If $e = 1$ the firm obtains 0 regardless of the state therefore $s_{pp} = b_i(S^*) = S^* > s^*$.

tation. To provide experimentation incentives, the planner needs to be more lenient in adoption $S_{pp} < S^*$ or/and more forgiving in withdrawal $z_{pp} < z^*$. As shown in Proposition 6, whenever $e > \hat{e}$, the planner is more lenient with both standards since in both cases the second-order losses associated with deviations from the optimum are swamped by the first-order gains in experimentation incentives.

Figure 7 plots the standards s_{pp}, z_{pp}, S_{pp} as a function of the externality, and compares them to standards at the social otpimum. For $e \leq \hat{e}$, all three standards are at the socially optimal level. For $e > \hat{e}$, as e increases, the rejection standard s^*_{pp} increases, and in response both the adoption and withdrawal standards decrease.

6 Private Information Collection and Costly Lying

In the previous sections, information was publicly disseminated. However firms, being in control of the research process before approval, have private access to research results and could potentially lie about them. For example, pharmaceutical firms may misrepresent the evidence presented to the FDA, the medical profession, or the public at large. The alleged withholding of negative results by pharmaceutical companies in the recent cases of Vioxx (an anti-inflammatory drug proven to increase the risk of cardiovascular events) or Paxil (an anti-depressant that could increase the suicide rates among children) generated major uproar and large demands for compensation.

We now depart from the public information assumption and assume that information obtained during the first experimentation phase is privately observed and non verifiable, so that the firm can make any report regardless of actual knowledge held by the firm. However, we assume that misreporting is costly.[34] Misreporting involves in practice risks both in terms of reputation as well as financial sanctions. Moreover, the probability that a lie is detected or condemned increases in the size of the lie. For instance, in the case of Vioxx, it is because Merck was shown to have withheld evidence that the penalties were very large: it allegedly paid over 4.85 billion dollars for settling individual complaints from patients.

Specifically we assume that, if the state is bad and the firm, when approval is granted, has collected evidence showing that the belief is S_i and chooses to report S, the probability of obtaining a given fine F is given by $P(S - S_i)$ where P is increasing and convex and $P(0) = P'(0) = 0$.

We consider F as an additional instrument and show that, when the planner cannot charge liability, costly lying will actually be beneficial and the fine will be set at intermediate levels that allow for some lying in equilibrium. In Appendix A.1, we consider fixed values of F and show

[34] We build on Kartik, Ottaviani, and Squintani's (2007) specification of costly lying.

that the results of Proposition 6 are preserved under costly lying for F sufficiently high.

6.1 Tolerating Lies

In many instances, the use of ex ante regulation shields the firm from product liability. As explained in Section 4 this is the case in the US for generic drugs and medical devices, where the approval by the FDA shields from product liability, but not from suits alleging misreporting. This is also the case more generally in other developed countries, in particular in Europe.

So we now consider a case where information is privately collected by the firm, the planner does not have access to liability and can commit at date $t = 0$ to a required report for approval, a withdrawal standard and a fine for misreporting (S, z, F). The following result shows that the fine, for low levels of externality, will be set at intermediate levels that tolerate some lying in equilibrium. Furthermore the planner is better off than if information was publicly collected.

Proposition 7 **(a)** *In an optimal mix of ex ante, ex post commitments and fines for misreporting* (S, z, F) *there exists* $\hat{e}$ *such that:*

- **(i)** *The first best is achieved if and only if* $e < \hat{e}$ *and all three instruments are used,.*
- **(ii)** *If* $e < \hat{e}$*, the optimal fine is interior, first decreasing then increasing and the firm lies in equilibrium:* $S_i < S$*. If* $e > \hat{e}$*, the optimal fine is* $F = \infty$ *and the firm does not lie.*

(b) *In an optimal mix of instruments, the* ***ex post commitment is chosen equal or more lenient*** *than the socially optimal levels* z^* *and the belief of the firm at approval* S_i ***is lower or equal*** *than the socially optimal approval standard* S^**.*

(c) *Welfare under costly lying is higher than welfare under public information.*

Whenever an interior fine is chosen, the firm lies in equilibrium. This is directly implied by the fact that the marginal cost of lying at zero is zero ($P'(0) = 0$). However, the planner knows precisely the size of the lie in equilibrium. Indeed for any approval standard S required for approval, even though the firm can make any report, the planner knows at what level S_i the firm stopped searching given that the firm takes into account the expected level of penalty it will incur for any lie $S - S_i$. Our planner optimally selects a combination of fine and report (F, S) that satisfies the firm's first order conditions (4) and (5) in Appendix A.[35]

Figure 8 represents the firm's best replies for two levels of required reports (S' and S^*), when the externality is fixed at $e' < \hat{e}$. If the required report is $S = S'$, then the firm's best replies intersect at the socially optimal level S^*. If, instead, the planner requests a report $S = S^*$, keeping the other tools unchanged, the lower best reply shifts to the left, since for any given belief S_i at

[35] The planner maximizes its expected payoff for (S, S_i, s_i, F, z) subject to (4) and (5) à la Stackelberg.

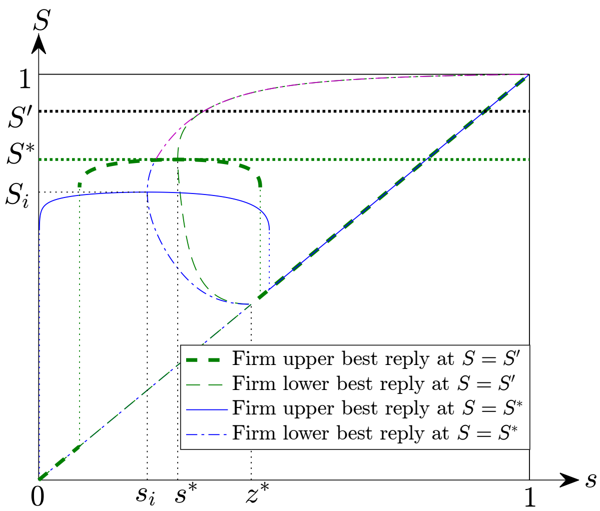

Figure 8: Best replies under costly lying for $e = e'$.

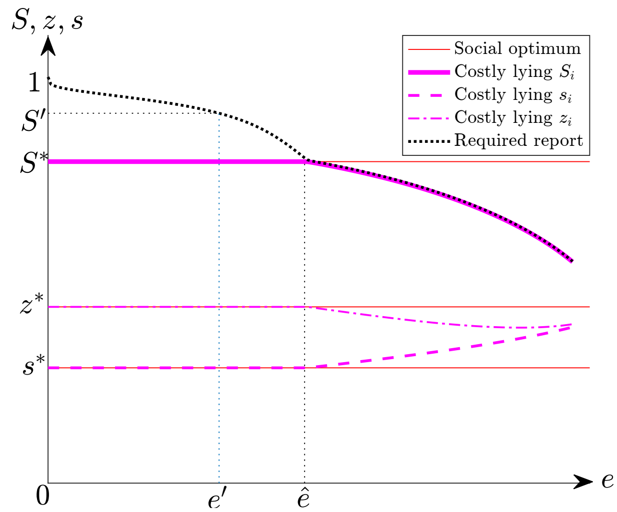

Figure 9: Optimal standards under costly lying as a function of e.

which approval is requested, the size of the lie is smaller ($S^* - S_i < S^{'} - S_i$) and thus the expected payoff in the bad state is higher, encouraging the firm to reject later. Overall this leads the firm to adopt at a lower level denoted S_i in Figure 8.

In the absence of liability, the fine for misreporting F serves two purposes. First, as explained above, for a given required report S, the choice of the fine determines the actual belief S_i at which the firm approves. Second, the fine, as a substitute for liability, serves to decrease excessive incentives for research when e is low.

Part (a-ii) of Proposition 7 describe the behavior of the optimal fine as a function of the externality. The fine F in combination with the required report to obtain approval S define the actual knowledge at approval S_i. When $e = 0$ the planner selects a report standard $S = 1$ forcing the probability of being caught lying to 1. In this way, the fine F serves exactly as a liability and the planner can perfectly align firm incentives in the bad state. As the externality increases, the planner cannot simply rely on the fine F to align incentives. Keeping the required report S high means it is too costly for the firm to influence and reduce the size of the lie.[36] Thus, the planner optimally decreases both the fine level and the required report standard S. This way, the firm can effectively influence the size of the lie reducing the expected penalty in the bad state and the first best obtains for low externality.

The optimal standards (S, S_i, s, z) as a function of e are plotted in Figure 9. As explained above, when the externality is low, the withdrawal standard is fixed at the socially optimal value. The required report is chosen above the socially optimal approval standard, but in combination

[36] The resulting regime is equivalent to a situation in which the planner imposes both liability and withdrawal commitments on the firm.

with an appropriately chosen fine, induces the firm to ask for approval at S^*. For instance, when the externality is $e^{'}$, the required report is set at $S^{'}$, and the firm submits at S^*, as we had illustrated in Figure 8. Finally, the fine is precisely chosen in a way to induce the correct research incentive. However, in the spirit of Proposition 6, when the externality is sufficiently high, the planner gives up on aligning the incentives of the firm by imposing liability in the bad state because the liability would reduce too much the firm's ex ante incentives for experimentation. This logic remains valid also when information is private. It is then optimal for the planner to set the fine as high as possible in order to kill the firm's incentives to lie. This way, no lying results in equilibrium and the firm adopts exactly at the required report level $S_i = S$, shutting down any risk of being caught lying and of incurring any penalty. In equilibrium, the planner provides research incentives to the firm by being more lenient both in the required report S and in the withdrawal standard z, analogously to the public information environment presented in Section 5.

To sum up, for low externality $e < \hat{e}$ the firm lies in equilibrium, the fine level is interior and the first best is achieved. For $e > \hat{e}$ there is no lying in equilibrium and the planner relies only on ex ante regulation through the required report standard S and on ex post regulation to incentivize experimentation. Overall in Proposition 7-c, when the planner is prevented to use liability, private information collection by the firm can be welfare improving. Because of costly lying, the planner knows the size of the lie in equilibrium and the fine for lying thus gives rise to an additional instrument, substitute for liability, that serves to chill research when it is socially desirable to do so, i.e., when the externality is small.

7 Conclusion

The paper formulates a tractable model to study the optimal mix of ex ante approval regulation, ex post withdrawal regulation, and liability to manage the adoption of activities with an externality of uncertain sign. The model combines optimal experimentation before implementation of the activity with learning after implementation. According to our main result, for low externality levels, using all three tools achieves the first best; when instead the externality is sufficiently high, liability would chill experimentation incentives too much, and in equilibrium a more lenient combination of only ex ante and ex post regulation is optimal. For the baseline model we focused on the case with public information, and then extended the analysis to the case with privately collected information that can be misreported.

Our model has a number of main applications discussed in detail in Section 4, but can be applied much more broadly as it for instance also speaks to other public policies that can be tested before approval (for instance through randomized or gradual implementation) and that then can be withdrawn following unsatisfactory results. We show that the predictions of the model are consistent with a number of regularities we identify. First, regulatory commitments in practice, and in particular ex post commitments, are typically designed to be weak. Second, regulation tends to be preferred to liability when the externality is large. Third, litigation and ex post regulation are rarely used together.

To conclude, let us mention one regularity not previously discussed. Most applications are characterized by the use of essentially a single instrument, the main exception being patent regulation that combines ex ante and ex post regulations. The use of limited instruments is typical for product safety, which mostly focuses on liability. This exclusive use of liability is somewhat at odds with the prediction of Proposition 6 that for low levels of the externality, all three instruments should be used in an optimal mix. When e is small, if the planner is constrained not to use regulation, the welfare loss is minimal since at $e = 0$ all standards are at the socially optimal level and moving away from them only induces second order losses. Moreover, it is reasonable to think that implementing an ex ante regulatory system and/or an ex post system of surveillance entails some costs. In this case, even small implementation costs would imply that only liability is used for low externality levels. Explaining this regularity could be the object of interesting future work.

8 References

Arkin, V. I., and Slastnikov, A. D. (2013). "On the Threshold Strategies in Optimal Stopping Problems for Diffusion Processes." *Journal of Applied Probability*, 54(3), 963–969

Arrow, K. J., and Fisher, A. C. (1974). "Environmental Preservation, Uncertainty, and Irreversibility." *Quarterly Journal of Economics*, 88(2): 312–319.

Bergemann, D. and Välimäki, J. (2008). "Bandit Problems." In: Durlauf, S. N., and Blume, L. E. (Eds.), *The New Palgrave Dictionary of Economics*, Second Edition, Palgrave Macmillan.

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

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

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

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

Chan, J., Lizzeri, A., Suen, W., and Yariv, L. (2018). "Deliberating Collective Decisions." *Review of Economic Studies*, 85(2): 929–963.

Che, Y.-K., and Mierendorff, K. (2017). "Optimal Sequential Decision with Limited Attention." Columbia and UCL Working Paper.

Coyle, P. (2012). "Uniform Patent Litigation in the European Union: An Analysis of the Viability of Recent Proposals Aimed at Unifying the European Patent Litigation System." *Washington University Global Studies Law Review*, 11(1): 171–192.

Daemmrich, Arthur (2007). "Pharmacovigilance and the Missing Denominator: The Changing Context of Pharmaceutical Risk Mitigation." *Pharmacy in History*, 49(2): 61–75.

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

Dixit, A. K., and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press.

Dvoretsky, A., Keifer, J., and Wolfowitz, J. (1953). "Sequential Decision Problems for Processes with Continuous Time Parameter: Testing Hypothesis." *Annals of Mathematical Statistics*, 24(2): 254–264.

Ford, R. (2013). "Patent Invalidity versus Noninfringement." *Cornell Law Review*, 99(1): 71–129.

Friedman, E., and Wickelgren, A. L. (2017). "Who (if Anyone) Should be Liable for Injuries from Generic Drugs?" *Journal of Law, Economics, and Organization*, 33(3): 541–577.

Friehe, T. and Schulte, E. (2017). "Uncertain product risk, information acquisition and product liability." *Economics Letters*, 159: 92–95.

Fudenberg, D., Strack, P., and Strzalecki, T. (2018). "Speed, Accuracy, and the Optimal Timing of Choices." *American Economic Review*, forthcoming.

Galasso, A., and Schankerman, M. (2015). "Patents and Cumulative Innovation: Causal Evidence from the Courts." *Quarterly Journal of Economics*, 130(1): 317–369.

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

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

Grenadier, S. R., Malenko, A., and Malenko, N. 2016. “Timing Decisions in Organizations: Communication and Authority in a Dynamic Environment.” *American Economic Review*, 106(9): 2552–2581.

Gul, F., and Pesendorfer, W. (2012). “The War of Information.” *Review of Economic Studies*, 79(2): 707–734.

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

Halac, M., Kartik, N., and Liu, Q. (2016). “Optimal Contracts for Experimentation.” *Review of Economic Studies*, 83(3): 1040–1091.

Hall, B. H., Graham, S., Harhoff, D., and Mowery D. C. (2004). “Prospects for Improving U.S. Patent Quality via Postgrant Opposition.” *Innovation Policy and the Economy*, 4: 115–143.

Henderson, J. A. Jr., and Twerski, A. D. (1998). “Achieving Consensus on Defective Product Design.” *Cornell Law Review*, 83(4): 867–920.

Henry, C. (1974). “Investment Decisions under Uncertainty: The Irreversibility Effect.” *American Economic Review*, 64(6): 1006–1012.

Henry, E., and Ottaviani, M. (forthcoming). “Research and The Approval Process: The Organization of Persuasion.” *American Economic Review*.

Hua, X., and Spier, K. E. (2018). “Product Safety, Contracts, and Liability.” SSRN Working Paper.

Jaffe, B., and 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.

Kartik, N., Ottaviani, M., and Squintani, F. (2007). “Credulity, Lies, and Costly Talk.” *Journal of Economics Theory*, 134(1): 93–116.

Ke, T. T., Shen, Z.-J. M., and Villas-Boas, J. M. (2016). “Search for Information on Multiple Products.” *Management Science*, 62(12): 3576–3603.

Ke, T. T., and Villas-Boas, J. M. (2018). “Optimal Learning before Choice.” MIT and Berkeley Working Paper.

Kolstad, C. D., Ulen, T. S., and Johnson, G. V. (1990). "Ex Post Liability for Harm vs. Ex Ante Safety Regulation: Substitutes or Complements?" *American Economic Review*, 80(4): 888–901.

Lybecker, K. M., and Watkins, L. (2015). "Liability Risk in the Pharmaceutical Industry: Tort Law in the US and UK." *Social Science Journal*, 52(4): 433–448.

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

McClellan, A. (2017). "Experimentation and Approval Mechanism." New York University Working Paper.

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

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

Orlov, D., Skrzypacz, A., and Zryumov, P. (2018). "Persuading the Principal to Wait." SSRN Working Paper.

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

Ottaviani, M., and Wickelgren, A. L. (2011). "Ex Ante or Ex Post Competition Policy? A Progress Report." *International Journal of Industrial Organization*, 29(2): 356–359.

Peskir, G., and Shiryaev, A. N. (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhäuser, Lecture Notes in Mathematics, ETH Zurich.

Posner, E. A., and Weyl, E. G. (2013). "An FDA for Financial Innovation: Applying the Insurable Interest Doctrine to 21st Century Financial Markets." *Northwestern University Law Review*, 107(3): 1307–1358.

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

Rupp, N. G., and Taylor, C. R. (2002). "Who Initiates Recalls and Who Cares? Evidence from the Automobile Industry." *Journal of Industrial Economics*, 50(2): 123-149.

Schmitz, P. W. (2000). "On the Joint Use of Liability and Safety Regulation." *International Review of Law and Economics*, 20(3): 371–382.

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

Shavell, S. (1984). "Liability for Harm Versus Regulation of Safety." *Journal of Legal Studies*, 13(2): 357–374.

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

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

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.

Viscusi, W. K., Vernon, J. M., and Harrington, J. E. (1995). *Economics of Regulation and Antitrust*. MIT Press.

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

Welling, L. (1991). "A Theory of Voluntary Recalls and Product Liability." *Southern Economic Journal,* 57(4): 1092–1111.

Zuckerman, D. M., Brown, P., and Nissen, S. E. (2011). "Medical Device Recalls and the FDA Approval Process." *Archives of Internal Medicine*, 171(11): 1006–1011.

A Appendix A: Derivations and Proofs

The two general lemmas that follow play a key role in our analysis.

Lemma 1 *In the Bandit problem in the ex post phase, if it is initially optimal to withdraw W, in the sense that there exists $\delta > 0$ such that $\hat{u}_p^2(\sigma_0) \equiv \sup_{(d_t)_{t\geq 0}} u_p^2(\sigma_0, d) = \left\{u_p^2(\sigma_0) | d_t = 0, 0 \leq t \leq \delta\right\}$, then it is always optimal to choose W.* $d = (d_t)_{t\geq 0}$ is the decision rule, $d_t \in \{0,1\}$ respectively for actions W and C,

Proof of Lemma 1

Let the payoff of the planner in the ex post phase be given by:

$$\hat{u}_p^2(\sigma_0) \equiv \sup_d \mathbb{E}_{\Omega\times\Theta}\left[\int_0^\infty e^{-rt} d_t \left[q_t v_p^G + (1-q_t)\, v_p^B\right] dt \,\middle|\, \sigma_0\right],$$

Fix $\varepsilon > 0$ and $d = (d_t)_{t\geq 0}$ such that $d_t = 0 \;\forall\, 0 \leq t \leq \delta$ and $u_p^2(\sigma_0, d) \geq \hat{u}_p^2(\sigma_0) - \varepsilon$. Then we have

$$\begin{aligned} u_p^2(\sigma_0, d) &= \mathbb{E}_{\Omega\times\Theta}\left[\int_\delta^\infty e^{-rt} d_t [q_t v_p^G + (1-q_t) v_p^B] dt \,\middle|\, \sigma_0\right] \\ &= e^{-r\delta}\mathbb{E}_{\Omega\times\Theta}\left[\int_0^\infty e^{-rt'} d_{t'} [q_{t'} v_p^G + (1-q_{t'}) v_p^B] dt' \,\middle|\, \sigma_0'\right] \leq e^{-r\delta}\hat{u}_p^2(\sigma_0), \end{aligned}$$

where $t' = t - \delta$ and $q_{t'} \equiv q_{t+\delta}$. Therefore, we conclude that $\hat{u}_p^2(\sigma_0) - \varepsilon \leq e^{-r\delta}\hat{u}_p^2(\sigma_0)$ from which $\hat{u}_p^2(\sigma_0) \leq \frac{\varepsilon}{1-e^{-r\delta}}$. Since ε was arbitrary then $\hat{u}_p^2(\sigma_0) \leq 0$ which is achievable by always pulling the safe arm W.

Lemma 2 *Let X_t be a one-dimensional stochastic process that solves $dX_t = \mu(X_t)dt + \sigma(X_t)dW_t$, where W_t is a Wiener process. Then for the problem*

$$\sup_{(\tau, d_\tau)} \mathbb{E}\left[e^{-r\tau}\left(d_\tau g_1(X_\tau) + (1-d_\tau) g_2(X_\tau)\right) \middle| X_0 = x\right]$$

there exists a solution of the form $\tau = \inf\{t : X_t \notin (\underline{x}, \bar{x})\}$ with $d_\tau = \mathbf{1}(X_\tau = x)$ for $x = \underline{x}$ and $x = \bar{x}$.

Proof of Lemma 2

Given any stopping time τ then $d_\tau = 1 \iff g_1(X_\tau) \geq g_2(X_\tau)$. Therefore defining $g(X_\tau) = \max\{g_1(X_\tau), g_2(X_\tau)\}$ we can rewrite this as an optimal stopping problem,

$$\sup_\tau \mathbb{E}\left[e^{-r\tau} g\left(X_\tau\right) | x\right].$$

It follows from Peskir and Shiryaev (2006) that an optimal stopping time is the first time at which X_t exits the continuation set $C = \{x : U(x) > g(x)\}$ where $U(x) \equiv \sup_\tau \mathbb{E}\left[e^{-r\tau} g(X_\tau)|x\right]$ is the value function. Under our assumptions C is an open set in R and therefore it can be represented as a countable union of disjoint (open) intervals. For this reason, the problem of determining the optimal stopping time can be reduced to determining an optimal first exit time from an interval $(\underline{x}, \overline{x})$ which contains the initial point x of the process X_t. In general, even if the threshold strategy is optimal, the continuation set C may not have a threshold structure. For instance, let $C = \{x : U(x) > g(x)\} = (x_1, x_2) \cup (x_3, x_4)$ with $x_1 \leq x_2 \leq x_3 \leq x_4$ and $D = \{x : U(x) = g(x)\}$ the stopping region. Set C does not have a threshold structure but given $x \in C$ being the process continuous we can find $C' = (\underline{x}', \overline{x}')$ which achieves the same expected payoff when starting at x such that $\underline{x}' = \sup_y \{y \in \partial C : y \leq x\}$ and $\overline{x}' = \inf_y \{y \in \partial C : y \geq x\}$. Hence, there is always an optimal stopping policy in the form of a threshold strategy around $X_0 = x$.

Proof of Proposition 1

The model is solved by backward induction starting from the ex post monitoring phase.

Step 1: *The optimal solution to the ex post problem requires the planner to select the withdrawal threshold z^* independent from the belief at which the agent chooses to adopt.*

Let S be the belief at which the social planner adopts. By Lemma 1, the bandit problem in the ex post phase can be rewritten as the stopping problem

$$\hat{u}_p^2(\mathsf{S}) \equiv \sup_{\tau} \mathbb{E}_{\Omega \times \Theta} \left[(1 - e^{-r\tau}) \left(q_\tau \frac{v_p^G}{r} + (1 - q_\tau) \frac{v_p^B}{r} \right) \middle| \mathsf{S} \right].$$

By Lemma 2 we know that there exists an optimal policy with two thresholds, denoted by (z, Z), around the initial belief S. Following Stokey (2009, Chapter 6), this ex post monitoring phase setup is a one-time one-option problem in which the agent selects the withdrawal (lower) threshold z below which it is optimal to stop the flow of benefits and withdraw, whereas the higher threshold is $Z = 1$. The planner's ex post utility is then

$$u_p^2(\mathsf{S}, \tau(z)) = \frac{e^{\mathsf{S}}}{1 + e^{\mathsf{S}}} \left(\frac{v_p^G}{r} \right) \left(1 - \mathbb{E}_\Omega \left[e^{-r\tau} | G, S \right] \right) + \frac{1}{1 + e^{\mathsf{S}}} \left(\frac{v_p^B}{r} \right) \left(1 - \mathbb{E}_\Omega \left[e^{-r\tau} | B, S \right] \right).$$

As in Stokey (2009), we have a closed-form expression for $\mathbb{E}_\Omega [e^{-r\tau} | \theta, \sigma] = \psi(\sigma, \theta, z)$ with $\psi(\sigma, G, z) = e^{-r_2(\sigma - z)}$ and $\psi(\sigma, B, z) = e^{r_1(\sigma - z)}$, where $r_1 = \frac{1}{2} \left(1 - \sqrt{1 + \frac{4r}{\mu_2'}} \right) < 0$ and $r_2 = \frac{1}{2} \left(1 + \sqrt{1 + \frac{4r}{\mu_2'}} \right) > 0$, $r_1 < r_2$, and $r_1 + r_2 = 1$. Therefore, the planner's problem becomes

$$\max_z \left\{ \frac{e^{\mathsf{S}}}{1 + e^{\mathsf{S}}} \left(\frac{v_p^G}{r} \right) \left(1 - e^{-r_2(\mathsf{S} - \mathsf{z})} \right) + \frac{1}{1 + e^{\mathsf{S}}} \left(\frac{v_p^B}{r} \right) \left(1 - e^{r_1(\mathsf{S} - \mathsf{z})} \right) \right\},$$

with solution $\mathsf{z}^* = -\ln \frac{v_p^G r_2}{v_p^B r_1}$, which is independent from the initial adoption belief S.

Step 2: *In the ex ante experimentation phase the social planner selects two thresholds (s^*, S^*) independent from the initial belief.*

By Step 1 we know that regardless of the belief at which the adoption decision is taken the planner always selects the same optimal threshold z^*. Thus, z^* pins down the ex ante expected payoff upon adoption

$$\hat{u}_p^2(\mathsf{S}) = \frac{e^{\mathsf{S}}}{1+e^{\mathsf{S}}}\left(\frac{v_p^G}{r}\right)\left(1-e^{-r_2(\mathsf{S}-\mathsf{z}^*)}\right) + \frac{1}{1+e^{\mathsf{S}}}\left(\frac{v_p^B}{r}\right)\left(1-e^{r_1(\mathsf{S}-\mathsf{z}^*)}\right),$$

where S is the belief at which the planner adopts. The planner then solves the following Wald problem,

$$\hat{u}_p^1(\sigma_0) \equiv \sup_{\tau,d_\tau} \mathbb{E}_{\Omega\times\Theta}\left[d_\tau\left(e^{-r\tau}\hat{u}_p^2(\sigma_\tau)\right) - \int_0^\tau e^{-rt}cdt \,\middle|\, \sigma_0\right].$$

Notice that for a given stopping time τ we have $d_\tau = 1 \iff \hat{u}_p(\sigma_\tau) \geq 0$. Therefore this is an optimal stopping with

$$\hat{u}_p^1(\sigma_0) \equiv \sup_{\tau} \mathbb{E}_{\Omega\times\Theta}\left[e^{-r\tau}\left(\hat{u}_p^2(\sigma_\tau)+\frac{c}{r}\right)\middle|\, \sigma_0\right] - \frac{c}{r}.$$

By Lemma 2 we know that the optimal solution consists in selecting a policy with two thresholds (s,S) around the initial belief σ_0, one for rejection and one for adoption. Following Stokey (2009, Chapter 5), define the expected discounted probabilities that given state θ the adoption and rejection standards are reached first respectively,

$$\begin{aligned}
\Psi(\sigma_0,\theta,s,S) &= \mathbb{E}_\Omega[e^{-r\tau}|\sigma_\tau = \mathsf{S},\theta,\sigma_0]\Pr(\sigma_\tau = \mathsf{S}|\theta,\sigma_0) \\
\psi(\sigma_0,\theta,s,S) &= \mathbb{E}_\Omega[e^{-r\tau}|\sigma_\tau = \mathsf{s},\theta,\sigma_0]\Pr(\sigma_\tau = \mathsf{s}|\theta,\sigma_0)
\end{aligned}$$

where $\Pr(A) = P_t\{\sigma_t^{-1}(\omega,\theta)\} = P_t\{(\omega,\theta)\in\Omega\times\Theta | \sigma_t(\omega,\theta)\in A\}$ and A is a Borel set. Hence we can rewrite the planner's utility as

$$u_p^1(\sigma_{0,}\tau(s)\wedge\tau(S)) = \frac{e^{\sigma_0}}{1+e^{\sigma_0}}u^1(\sigma_{0,}\tau(s)\wedge\tau(S),G) + \frac{1}{1+e^{\sigma_0}}u_p^1(\sigma_{0,}\tau(s)\wedge\tau(S),B),$$

where for $\theta\in\{A,B\}$

$$\begin{aligned}
u_p^1(\sigma_{0,}\tau(s)\wedge\tau(S),\theta) &= \sum_{\gamma\in\{\mathsf{s},\mathsf{S}\}}\Pr(\sigma_\tau=\gamma|\theta,\sigma_0)\mathbb{E}_\Omega\left[e^{-r\tau}\hat{u}_p^2(\sigma_\tau,\theta) - \int_0^\tau ce^{-rt}dt\,\middle|\,\sigma_\tau=\gamma,\sigma_0,\theta\right] \\
&= -\frac{c}{r} + \left(\hat{u}_p^2(\mathsf{S},\theta)+\frac{c}{r}\right)\Psi(\sigma_0,\theta,\mathsf{s},\mathsf{S}) + \frac{c}{r}\psi(\sigma_0,\theta,\mathsf{s},\mathsf{S})
\end{aligned}$$

and $\hat{u}_p^2(\sigma_\tau,\theta) = \left(\frac{v_p^G}{r}\right)(1-\psi(\sigma_\tau,\theta,\mathsf{z}^*))$. The planner solves $\max_{\mathsf{s},\mathsf{S}} u_p^1(\sigma_{0,}\tau(s)\wedge\tau(S))$. Using the closed-form expressions for Ψ and ψ given in Henry and Ottaviani (forthcoming),[37] the first order conditions are

$$\frac{\partial u_p^1}{\partial s} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}}\psi\left\{a\left[u_p^2(S,G)+e^{-S}u_p^2(S,B)+\left(1+e^{-S}\right)\frac{c}{r}\right]+b\left(1+e^{-s}\right)\frac{c}{r}-e^{-s}\frac{c}{r}\right\} \tag{2}$$

[37] We have $\Psi(\sigma,G,\mathsf{s},\mathsf{S}) = \frac{e^{-R_1(\sigma-\mathsf{s})}-e^{-R_2(\sigma-\mathsf{s})}}{e^{-R_1(\mathsf{S}-\mathsf{s})}-e^{-R_2(\mathsf{S}-\mathsf{s})}}$ and $\psi(\sigma,G,\mathsf{s},\mathsf{S}) = \frac{e^{R_2(\mathsf{S}-\sigma)}-e^{R_1(\mathsf{S}-\sigma)}}{e^{R_2(\mathsf{S}-\mathsf{s})}-e^{R_1(\mathsf{S}-\mathsf{s})}}$; see Appendix A of Henry and Ottaviani (2017) for a detailed derivation. Moreover, by their Lemma B0 $\Psi(\sigma,B,\mathsf{s},\mathsf{S}) = e^{\sigma-\mathsf{S}}\Psi(\sigma,G,\mathsf{s},\mathsf{S})$ and $\psi(\sigma,B,\mathsf{s},\mathsf{S}) = e^{\sigma-\mathsf{s}}\psi(\sigma,G,\mathsf{s},\mathsf{S})$.

$$\frac{\partial u_p^1}{\partial S}=\frac{e^{\sigma_0}}{1+e^{\sigma_0}}\Psi\left\{f\left[u_p^2(S,G)+e^{-S}u_p^2(S,B)+\left(1+e^{-S}\right)\frac{c}{r}\right]+g\left(1+e^{-s}\right)\frac{c}{r}-e^{-s}\left[\frac{c}{r}+u_p^2(S,B)\right]\right\}, \tag{3}$$

from which we can verify that the optimal solution s^* and S^* are independent from the initial belief σ.[38]

Step 3: *The optimal triplet solution (s^*, S^*, z^*) is such that $s^* \leq z^* \leq S^*$.*
The continuation value with adoption at belief σ is given by $\hat{u}_p^2(\sigma)$. This is always nonnegative, $\hat{u}_p^2(\sigma) \geq 0$. In particular, we have that $\hat{u}_p^2(\sigma) = 0$ for all $\sigma \leq \mathrm{z}^*$ and $\hat{u}_p^2(\sigma) > 0$ and strictly increasing for $\sigma > z^*$. Therefore, the option value of experimenting is non-negative around z^*, the highest belief at which the planner is indifferent between rejecting and adopting. Step 2 establishes result (a), Step 1 proves result (b) and Steps 1-3 are used to prove result (c).

Proof of Proposition 2
Step 1: *Adoption and rejection standards are decreasing in v_p^G.*
By the implicit function theorem it is enough to study the sign of $\left.\frac{\partial u_p^1}{\partial s \partial v_p^G}\right|_{s=b_p(S)}$ to infer the sign of $\frac{\partial b_p(S)}{\partial v_p^G}$. Taking the derivative we have

$$\frac{\partial u_p^1}{\partial s \partial v_p^G}=\frac{e^{\sigma_0}}{1+e^{\sigma_0}}\psi a\left\{\left[u_{p,z}^2(S,G,z)+e^{-S}u_{p,z}^2(S,B,z)\right]\frac{\partial z}{\partial v_p^G}+\frac{\left(1-e^{-r_2(S-z)}\right)}{r}\right\}.$$

By the Envelope Theorem the first term in the square brackets is zero, so that

$$\frac{\partial u_p^1}{\partial s \partial v_p^G}=\frac{e^{\sigma_0}}{1+e^{\sigma_0}}\psi a\frac{\left(1-e^{-r_2(S-z)}\right)}{r}<0,$$

which is negative since $a < 0$ and the fraction is clearly positive.

Following the same reasoning, we can show that $\frac{\partial B_p(s)}{\partial v_p^G} < 0$. Under strong liability, $L = \overline{L}$, the firm solves a stand-alone problem with insufficient incentives in the good state, $v_i < v_p^G$, hence $z_{ii}\left(\overline{L}\right) > z^*$, $s_{ii}\left(\overline{L}\right) > s^*$, and $S_{ii}\left(\overline{L}\right) > S^*$.

Step 2: *The optimal liability L^* is interior: $v_i < L^* < \overline{L}$.*

Notice that for any $L \leq v_i$ the firm as no incentive in conducting costly experimentation neither ex ante nor ex post. Indeed, $s_{ii} = z_{ii} = S_{ii} = 0$ since in both states the firm obtains a strictly positive payoff. Therefore it must be the case that $L^* > v_i$. From Step 1 we know that at $L = \overline{L}$ the firm incentives in the good state are insufficient; consequently the firm performs insufficient experimentation ex ante, adopts too late, and withdraws too early. By selecting $L < \overline{L}$

[38] The terms $a < 0, b > 0, f < 0, g > 0$ are functions of s and S and are defined in Appendix A Lemma B0 of Henry and Ottaviani (2017).

the planner increases the firm value in the bad state. Since the firm solves a stand alone problem, it is enough to study how the social planner reacts when v_p^B increases. The withdrawal standard decreases in v_p^B; from $z^* = -\ln \frac{v_p^G r_2}{v_p^B r1}$, we have $\frac{\partial z^*}{\partial v_p^B} = \frac{1}{v_p^B} < 0$. For the rejection standard, by the implicit function theorem it is enough to study the sign of $\left.\frac{\partial u_p^1}{\partial s \partial v_p^B}\right|_{s=b_p(S)}$. Taking the derivative and proceeding as in Step 1 we obtain

$$\frac{\partial u_{p,s}^1}{\partial v_p^B} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \psi a e^{-S} \frac{\left(1-e^{r_1(S-z)}\right)}{r} < 0.$$

For the adoption standard S^*, following the same reasoning we obtain $\frac{\partial u_{p1}}{\partial S \partial v_p^B} < 0$.

Therefore, by reducing the liability the planner pushes the firm to increase experimentation, to adopt earlier, and to withdraw later. Hence, it must be the case that $L^* < \overline{L}$.

Step 3: *There exists $\hat{L}$ such that $s^* \leq s_{ii}(\hat{L}) \leq S_{ii}(\hat{L}) \leq S^*$.*

We can easily solve for L such that $z^* = z_{ii}(L)$ to find that $\hat{L} = \frac{v_i(v_p^G - v_p^B)}{(v_p^G)}$. To show that the amount of ex ante experimentation under liability at $L = \hat{L}$ is lower than the socially optimal level, consider the single decision maker problem and perform comparative statics with respect to the changes in the payoffs v^G and v^B such that the withdrawal standard is kept at the same level. Recall that for a decision maker with payoffs $v^G > 0$ and $v^B < 0$, the optimal withdrawal standard is given by $z^* = -\ln\left(\frac{r_2 v^G}{r_1 v^B}\right)$. To keep z^* constant when we change the payoffs it must hold that

$$\frac{\partial v^B}{\partial v^G} = \frac{v^B}{v^G}.$$

The marginal value for any given S of rejecting earlier is given by

$$\left.\frac{\partial u^1(\sigma_0)}{\partial s}\right|_{s=b(S)} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \psi \left\{ a\left[u^2(S,G) + e^{-S} u^2(S,B) + \left(1+e^{-S}\right)\frac{c}{r}\right] + b\left(1+e^{-s}\right)\frac{c}{r} - e^{-s}\frac{c}{r} \right\}.$$

By the implicit function theorem, it suffices to study the sign of

$$\left.\frac{\partial u^1(\sigma_0)}{\partial s \partial v^G}\right|_{s=b(S)} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \psi \frac{a}{v^G}\left[u^2(S,G) + e^{-S} u^2(S,B)\right] < 0$$

for $S > z^*$, to conclude that $\frac{\partial b(S)}{\partial v^G} < 0$. Therefore, if v^G decreases and v^B increases to keep z^* unchanged, the ex ante rejection standard $s = b(S)$ is higher for any adoption standard $S > z^*$.

Turning to the marginal value of adopting later for any given s, we have

$$\left.\frac{\partial u^1(\sigma_0)}{\partial S}\right|_{S=B(s)} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \Psi \left\{ \begin{array}{c} f\left[u^2(S,G) + e^{-S} u^2(S,B) + \left(1+e^{-S}\right)\frac{c}{r}\right] + g\left(1+e^{-s}\right)\frac{c}{r} \\ -e^{-s}\left(\frac{c}{r} + u^2(S,B)\right) \end{array} \right\}.$$

Following the same reasoning, as before we have

$$\left.\frac{\partial u^1(\sigma_0)}{\partial S \partial v^G}\right|_{S=B(s)} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \Psi \frac{1}{v^G} \left\{ f\left[u^2(S,G) + e^{-S} u^2(S,B)\right] - e^{-S} u^2(S,B) \right\} > 0,$$

so that $\frac{\partial B(s)}{\partial v^G} > 0$. When v^G decreases and v^B increases to keep z^* unchanged, the adoption standard will be lower for any given s. Notice that the term $f\left[u^2(S,G) + e^{-S}u^2(S,B)\right] - e^{-S}u^2(S,B)$ is positive. Indeed at $S = B(s)$, we have

$$f\left[u^2(S,G) + e^{-S} u^2(S,B)\right] - e^{-S} u^2(S,B) = -\left[f\left(1+e^{-S}\right)\frac{c}{r} + g\left(1+e^{-s}\right)\frac{c}{r} - e^{-S}\frac{c}{r}\right].$$

The term in brackets on the right hand side must be negative, because it can be seen as the marginal value of approving later for an agent that obtains zero payoff regardless of the state and pays research cost $c > 0$.[39]

Step 4: *The optimal liability L^* is such that $s^* \le s_{ii}(L^*) \le S_{ii}(L^*) \le S^*$.*
If $L = \hat{L}$, the ex post value is maximized but the ex ante amount of experimentation is insufficient, as we know from Step 3. If $L^* \ge \hat{L}$ and $S_{ii}(L^*) \ge S^*$, reducing liability moves all standards closer to the socially optimal levels, hence it must be the case that $S_{ii}(L^*) \le S^*$. Analogously, we can conclude that $s^* \le s_{ii}(L^*)$.

Step 5: *The optimal liability L^* is a function of the initial belief q_0 and such that $L^* \le \hat{L}$ (resp. $L^* \ge \hat{L}$) if $q_0 \in (\underline{q}, \hat{q})$ (resp. $q_0 \in (\hat{q}, \overline{q})$).*
First recall that $u_p^1(q_0)$ is increasing in σ_0 and can be rewritten as

$$u_p^1(\sigma_{0,}, s, S) = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} u^1(\sigma_0, s, S, G) + \frac{1}{1+e^{\sigma_0}} u_p^1(\sigma_0, s, S, B),$$

where for $\theta \in \{G, B\}$

$$u_p^1(\sigma_0, s, S, \theta) = -\frac{c}{r} + \left[\hat{u}_p^2(S,\theta) + \frac{c}{r}\right] \Psi(\sigma_0, \theta) + \frac{c}{r}\psi(\sigma_0, \theta).$$

Clearly, if $q_0 \le s^*$ or $q_0 \ge S^*$, then the planner selects a liability level such that the firm immediately rejects in the first case or immediately adopts in the second case. Focusing on $q_0 \in (s^*, S^*)$, we have two cases. If $e = 0$, then $L^* = \hat{L} = \overline{L}$ for any q_0. If $e > 0$, we now show that it is optimal for the planner to impose a liability level $L^* \neq \hat{L}$. Suppose that $L = \hat{L}$, then $z_{ii}(\hat{L}) = z^*$; we now study the sign of

$$\frac{\partial u_p^1(\sigma_0)}{\partial L} = \frac{\partial u_p^1(\sigma_0)}{\partial z}\frac{\partial z}{\partial L} + \frac{\partial u_p^1(\sigma_0)}{\partial s}\frac{\partial s}{\partial L} + \frac{\partial u_p^1(\sigma_0)}{\partial S}\frac{\partial S}{\partial L}$$

evaluated at $(s_{ii}(\hat{L}), z_{ii}(\hat{L}), S_{ii}(\hat{L}))$. The first term is zero since $z_{ii}(\hat{L}) = z^*$, the second term is negative and the third is positive, hence the sign depends on the initial belief q_0. As $q_0 \to s_{ii}(\hat{L})$

[39] Given that zero is obtained regardless, it is optimal to set $B(s) = s$ for any given s. Thus, $\frac{\partial u^1(\sigma_0)}{\partial S} < 0$ for any s.

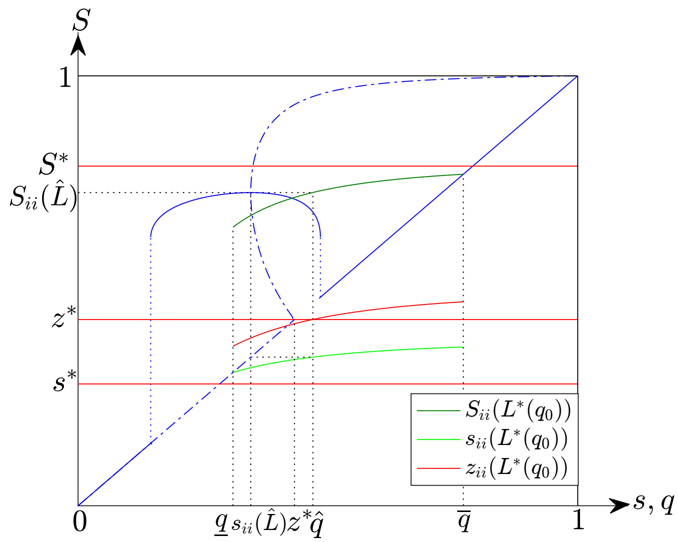

Figure 10: Policy paths as function of q_0 under liability.

the marginal value of delaying adoption goes to zero $\frac{\partial u_p^1(\sigma_0)}{\partial S} \to 0$ and $\frac{\partial u_p^1(\sigma_0)}{\partial L} = \frac{\partial u_p^1(\sigma_0)}{\partial s}\frac{\partial s}{\partial L} < 0$. Therefore, $L^* < \hat{L}$. By the same argument when $q_0 \to S_{ii}(\hat{L})$ we have $\frac{\partial u_p^1(\sigma_0)}{\partial L} = \frac{\partial u_p^1(\sigma_0)}{\partial S}\frac{\partial S}{\partial L} > 0$ and clearly, $L^* > \hat{L}$. By continuity there exists $\hat{q}$ such that $\frac{\partial u_p^1(\sigma_0)}{\partial s}\frac{\partial s}{\partial L} = -\frac{\partial u_p^1(\sigma_0)}{\partial S}\frac{\partial S}{\partial L}$ and $L^* = \hat{L}$.

As illustrated in Figure 10, $\underline{q} \in (s^*, \hat{q})$ and $\overline{q} \in (\hat{q}, S^*)$ are the beliefs at which the optimal liability L^* induces the firm to respectively reject immediately or adopt immediately. The three paths $(s_{ii}(L^*(q_0)), z_{ii}(L^*(q_0)), S_{ii}(L^*(q_0)))$ are increasing in the prior q_0 and as the externality $e \to 0$ these paths become flatter and converge (uniformly for any q_0) to the socially optimal standards (s^*, S^*, z^*).

Step 1-4 establishes result (a), Step 5 proves result (b).

Proof of Proposition 3

Step 1: *Adoption and rejection standards are increasing in z if and only if $z > z^*$.*

Let z be any exogenous withdrawal standard and let z^* be the optimal level. From the implicit function theorem and the concavity of the problem it is enough to focus on the term $\frac{\partial^2 u^1(\sigma)}{\partial S \partial z}$ to determine the sign of $\frac{\partial B(s)}{\partial z}$. Recall that

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

Taking the cross derivative with respect to z we obtain

$$\frac{\partial^2 u^1(\sigma)}{\partial S \partial z} = \frac{e^\sigma}{1+e^\sigma}\Psi\left\{ f\left(\frac{\partial u^2(S,G,z)}{\partial z} + e^{-S}\frac{\partial u^2(S,B,z)}{\partial z}\right) - e^{-S}\frac{\partial u^2(S,B,z)}{\partial z} + \frac{\partial u^2(S,G,z)}{\partial S \partial z} + e^{-S}\frac{\partial u^2(S,B,z)}{\partial S \partial z} \right\},$$

where $\frac{\partial u^2(S,G,z)}{\partial z} = -r_2(\frac{v^G}{r})e^{-r_2(S-z)} = -\frac{\partial u^2(S,G,z)}{\partial S}$ and $\frac{\partial u^2(S,Bz)}{\partial z} = r_1(\frac{v^B}{r})e^{r_1(S-z)} = -\frac{\partial u^2(S,B,z)}{\partial S}$. Noticing that $\frac{\partial u^2(S,G,z)}{\partial S \partial z} = r_2^2(\frac{v^G}{r})e^{-r_2(S-z)} = -r_2\frac{\partial u_2(S,G,z)}{\partial z}$ and $\frac{\partial u_2(S|B)}{\partial S \partial z} = r_1^2(\frac{v^B}{r})e^{r_1(S-z)} = r_1\frac{\partial u_2(S|B)}{\partial z}$, we have

$$\frac{\partial^2 u^1(\sigma)}{\partial S \partial z} = \frac{e^\sigma}{1+e^\sigma}\Psi\left\{ f\left(\frac{\partial u^2(S,G,z)}{\partial z} + e^{-S}\frac{\partial u^2(S,B,z)}{\partial z}\right) - e^{-S}\frac{\partial u^2(S,B,z)}{\partial z} - r_2\frac{\partial u^2(S,G,z)}{\partial z} + e^{-S} r_1 \frac{\partial u^2(S,B,z)}{\partial z} \right\}$$

and thus

$$\frac{\partial^2 u^1(\sigma)}{\partial S \partial z} = \frac{e^\sigma}{1+e^\sigma}\Psi(f-r_2)\left(\frac{\partial u^2(S,G,z)}{\partial z} + e^{-S}\frac{\partial u^2(S,B,z)}{\partial z}\right),$$

using $r_1 + r_2 = 1$. The sign of the numerator depends on the last term. Clearly, if $z > z^*$ the term is negative, so that the sign of $\frac{\partial B(s)}{\partial z}$ is positive. By the same argument, $\frac{\partial b(S)}{\partial z} > 0$ if and only if $z > z^*$. Notice that the optimal withdrawal for the firm would be $z_i^* = 0$; thus, whenever the social planner commits to a positive withdrawal, the firm increases both adoption and rejection standards.

Step 2: *There exists $\bar{e}$ such that $s_{ip}^* = s^*$ when the planner commits to $z_{ip} = z^*$.*
The rejection standard is decreasing in v_p^B when the planner commits to the optimal withdrawal z^*. Indeed, we have

$$\frac{\partial u_p^1}{\partial s \partial v_p^B} = \psi a e^{-S}\frac{\partial u_p^2(S,z^*,B)}{\partial v_p^B} < 0.$$

Hence, when $e = 0$ it must be the case $s_{ip} < s^*$ when the planner commits to the optimal ex post withdrawal. If, instead, $e \to 1$ the firm has no value in experimenting ex ante. Therefore, $s_{ip} \to z^* > s^*$. By continuity, there exists $\bar{e}$ such that $s_{ip} = s^*$ when the planner commits to the optimal withdrawal. With some algebra one can verify that

$$\frac{\partial u_i^1}{\partial s \partial e} = \frac{e^\sigma}{1+e^\sigma}\psi(\sigma,G,S,s)a\left[-\frac{v_p^G}{r}(1-\psi(S,G,z)) - e^{-S}\frac{v_p^G}{r}(1-\psi(S,B,z))\right] > 0,$$

so that $\frac{\partial s_{b_i}(S_{ip})}{\partial e} > 0$. Thus, $\bar{e}$ is also unique.

Step 3: *If $e \le \bar{e}$ then $z_{ip}^* > z^*$, whereas if $e > \bar{e}$ there exists $\bar{q}(e)$ such that if $q_0 \le \bar{q}(e)$ (resp. $q_0 > \bar{q}(e)$) then $z_{ip}^* < z^*$(resp. $z_{ip}^* > z^*$).*
If $e \le \bar{e}$ it is unambiguously optimal for our planner to commit to be tougher ex post $z_{ip}^* > z^*$. Indeed,

$$\frac{\partial u_p^1(\sigma_0)}{\partial z} = \frac{\partial u_p^1(\sigma_0)}{\partial S}\frac{\partial S}{\partial z} + \frac{\partial u_p^1(\sigma_0)}{\partial s}\frac{\partial s}{\partial z} + \frac{\partial u_p^1(\sigma_0)}{\partial z}$$

evaluated at (s_{ip}, z^*, S_{ip}) is positive. In fact, by Step 1 we know that the first and second term are positive, while the last term is zero since the withdrawal is set at the optimal level.

If $e > \bar{e}$ the role of the initial belief matters, since the second term of the expression above is now negative $\frac{\partial u_p^1(\sigma_0)}{\partial s}\frac{\partial s}{\partial z} < 0$. Overall the sign depends on the initial belief q_0. As $q_0 \to s_{ip}$, the marginal value of delaying adoption goes to zero. Hence, we are left with $\frac{\partial u_p^1(\sigma_0)}{\partial z} = \frac{\partial u_p^1(\sigma_0)}{\partial s}\frac{\partial s}{\partial z} < 0$, which implies that $z_{ip}^* < z^*$. The opposite holds if $q_0 \to S_{ip}$. In this case the planner commits to be tougher, $z_{ip}^* > z^*$. Therefore, there exists $\bar{q}(e) \in (s_{ip}, S_{ip})$ such that $z_{ip} = z^*$. This proves results (a) and (b) of the Proposition.

Proof of Proposition 4

Step 1: *Welfare under the liability regime is decreasing in e and tangent to the social optimum at $e = 0$.*

If $e = 0$, the optimal liability $L^* = \hat{L} = \overline{L}$ achieves the first best and the welfare is maximized. If $e > 0$ and the planner imposes L^* we know from the proof of Proposition 2 that $s^* \leq s_{ii}(L^*) \leq S_{ii}(L^*) \leq S^*$. Moreover the higher the externality level e, the lower the value of information and the lower the amount of experimentation in equilibrium, i.e., $S_{ii}(L^*) - s_{ii}(L^*)$ decreases in e. Therefore the welfare under optimal liability must be decreasing in e.

As soon as e moves away from zero, assuming the planner were to hold $L^*(0)$, both $s_{ii}(L^*)$ and $S_{ii}(L^*)$ would increase because the firm's expected payoff decreases. The planner can thus offset this effect by reducing the optimal liability since both $s_{ii}(L)$ and $S_{ii}(L)$ are increasing in L. This implies that around zero the welfare loss must be of second order. Hence, as in Figure 5, the welfare is tangent at $e = 0$ to the social optimum.

Step 2: *For e sufficiently low the welfare is higher under the liability regime than under ex-post regulation.*

From Proposition 3 we know that for $e < \overline{e}_{ip}$ the optimal ex post commitment to z_{ip}^* is necessarily tougher. Thus the welfare under ex post regulation is below the socially optimal level. From Step 1 above we know that welfare under liability is maximized for very low externality levels, thus for e sufficiently low we conclude that the liability regime welfare dominates.

Step 3: *For e sufficiently high and initial belief q_0 sufficiently low, ex-post regulation welfare dominates liability.*

Starting from the liability regime, a switch to an ex post regulation regime by keeping z unchanged and removing liability L result in a reduction of both rejection and adoption standards. Given withdrawal z, it is always the case that $s_{ii}(L(z)) \geq s_{ip}(z)$, so that ex post regulation reduces rejections. By the same argument as above, if e is sufficiently high and q_0 is sufficiently low, so that the relevant issue for the planner is reducing rejections, welfare is increased by moving from liability to ex post regulation, keeping the same withdrawal standard.

Proof of Proposition 5

Step 1: *The rejection standard is increasing in S.*

First, note that under the ex ante regulation the firm controls the withdrawal standard and thus sets $z_{pi} = 0$. The lower best reply is then implicitly defined for any given adoption standard S as

$$\frac{\partial u_i^1}{\partial s} = \frac{e^\sigma}{1+e^\sigma}\psi\left\{a\left(\frac{(1-e)v_p^G}{r} + \frac{c}{r}\right)(1+e^{-S}) + b(1+e^{-s})\frac{c}{r} - e^{-s}\frac{c}{r}\right\}.$$

By the implicit function theorem,

$$\frac{\partial u_i^1}{\partial s \partial S} = -a\psi\left(\frac{(1-e)v_p^G}{r} + \frac{c}{r}\right)e^{-S} > 0,$$

so that $b(S)$ is increasing in S.

Step 2: *There exists $\bar{e}$ such that $s_{pi} = s^*$ if the planner commits to the socially optimal adoption standard S^*.*

If $e = 0$, the firm's payoff is aligned to the planner's payoff in the good state, v_p^G. From Step 2 in Proposition 2 we know that $\frac{\partial u_p^1}{\partial s \partial v_p^B} < 0$. Therefore, it must be the case that $s_{pi} < s^*$. If $e \to 1$, the value of ex ante experimentation goes to zero, which implies that $s_{ip} \to S^* > s^*$, as in Figure 6. Thus, there exists an externality level $\bar{e}$ such that $s_{ip} = s^*$ which by Step 1 is also unique.

Step 3: *There exists $\check{e}$ and $\hat{e}$ such that $\bar{e} \in (\check{e}, \hat{e})$. Moreover if $e \in (\check{e}, \hat{e})$ the planner commits to be tougher $S_{pi}^* > S^*$. If $e \geq \hat{e}$ (resp. $e \leq \check{e}$) then there exists $\bar{q}(e)$ such that if $q_0 \leq \bar{q}(e)$ (resp. $q_0 \geq \bar{q}(e)$) the planner commits to be more lenient $S_{pi}^* < S^*$.*

Under Ex Ante regulation only, the planner has no way of change the ex post withdrawal incentives of the firm; thus, ex post withdrawal never takes place, $z_{pi} = 0$. The problem reduces to the Planner Commitment problem described by Henry and Ottaviani (forthcoming). Given that $z_{pi} = 0$, the social planner selects the adoption standard as in the classic Wald framework with only the ex ante phase.

Step 1 of Proposition 3 implies that when $z = 0$ the planner upper (resp. lower) best reply has to be above (resp. to the right) of the planner upper (resp. lower) best reply when $z = z^*$. Moreover, recall that as e increases the firm lower best reply $b_i(S)$ shifts to the right.[40]

Now define $\check{e}$ (resp. $\hat{e}$) be the externality level such that, when $z = 0$, the firm lower best reply $b_i(S)$ intersects the planner upper best reply $B_p(s)$ (resp. the planner lower $b_p(S)$) exactly at $(b_i(S^*), S^*)$. Moreover, for a given e, let $\tilde{S}$ be the adoption standard that is identified at the intersection between the lower best replies, where $b_i(\tilde{S}) = b_p(\tilde{S})$.[41]

Henry and Ottaviani (forthcoming) show that the planner commitment $S_{pi}^*(q_0) \in (\tilde{S}, S^N)$ is either increasing or decreasing in q_0.[42] In particular, if $e \in (\check{e}, \hat{e})$, then by construction $S_{ip}^*(q_0) > S^*$ for any q_0. If $e \geq \hat{e}$, the commitment path is increasing and $S^* \in (\tilde{S}, S^N)$; hence it is enough to define $\bar{q}(e)$ as the belief such that $S_{pi}^*(\bar{q}(e)) = S^*$. If instead $e \leq \check{e}$, the commitment path is decreasing and $S^* \in (S^N, \tilde{S})$; again, it is enough to define $\bar{q}(e)$ as the belief such that $S_{pi}^*(\bar{q}(e)) = S^*$.

[40] It is easy to verify that $\partial u_i^1 / \partial s \partial e > 0$.

[41] This point corresponds to the Informer Authority solution proposed in Henry and Ottaviani (2017) in the game between the planner and the firm.

[42] For given level of e, S^N identifies the adoption standard at the intersection between $b_i(S)$ and $B_p(S)$ when $z = 0$. These points correspond to the Nash solution in Henry and Ottaviani (2017).

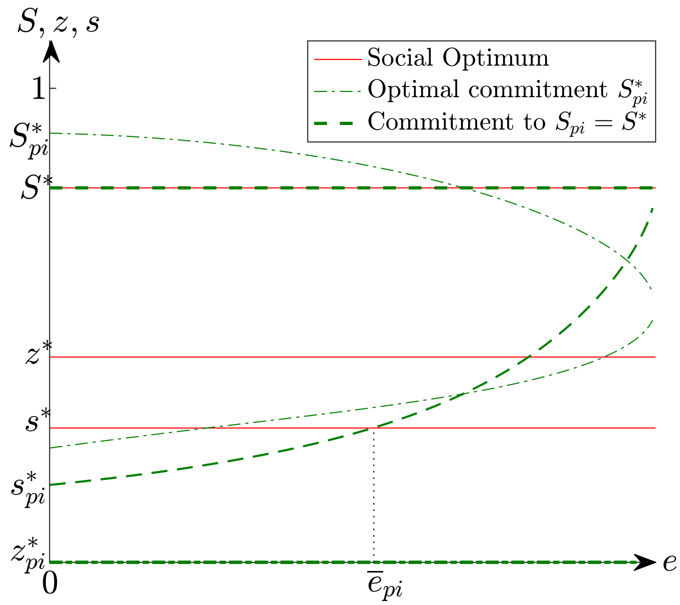

Figure 11: Optimal ex ante commitment.

We can thus conclude that if $e \geq \bar{e}$ there exists $\bar{q}(e) \geq 0$ such that whenever $q_0 \leq \bar{q}(e)$ the planner wants to be more lenient, $S^*_{pi} < S^*$. Figure 11 shows this property. The dotted green represents the optimal ex ante commitment solution as a function of the externality. The externality level $\hat{e}$ is such that $S^*_{pi}(\hat{e}) = S^*$, where the dashed-dotted green curve intersects the continuous horizontal line at S^*.

Proof of Proposition 6

Step 1: *There exists $\hat{e}$ such that the socially optimal policy (s^*, z^*, S^*) is achieved when the planner commits to $S_{pp} = S^*$ and $z_{pp} = z^*$.*

At $e = 0$, we know that the firm lower best reply is always to the left of the planner lower best reply. From $\frac{\partial u^1_p}{\partial s \partial v^B_p} < 0$, we have $s_{pp} < s^*$ when the planner commits to $S_{pp} = S^*$ and $z_{pp} = z^*$. If $e \to 1$ then the value of ex ante experimentation goes to zero, implying that $s_{pp} \to S^* > s^*$. By continuity there exists $\hat{e}$ such that $s_{pp} = s^*$; moreover, $\hat{e}$ is also unique since $\frac{\partial u^1_i}{\partial s \partial e} > 0$.

Step 2: *Suppose the planner can use all three instruments and denote the resulting policy as $(s_{pp}(L), z_{pp}(L), S_{pp}(L))$. If $e \leq \hat{e}$ the planner achieves the first best, by committing to $S_{pp} = S^*$ and $z_{pp} = z^*$.*

If $e = 0$ then strong liability $L = \overline{L}$ achieves the first best. If instead $e = \hat{e}$, Step 1 tells us that committing to the socially optimal withdrawal and adoption standards is enough to obtain the first best policy. Hence, in this case $L = 0$ and $s_{pp}(0) = s^*, S_{pp}(0) = S^*, z_{pp}(0) = z^*$.

Now fix $e \in (0, \hat{e})$. If $L = \overline{L}$ then $s^* < s_{pp}(\overline{L})$. To see this, notice that when $e = 0$ and $L = \overline{L}$ the planner and firm marginal value of rejecting earlier are perfectly aligned $\frac{\partial u^1_p}{\partial s} = \frac{\partial u^1_i}{\partial s}$, and $\frac{\partial u^1_i}{\partial s \partial e} > 0$.

If instead $L=0$ by Step 1 we have $s_{pp}(0) < s^*$. By continuity there exists $L(e) \in (0, \hat{L})$ such that $s_{pp} = s^*$; moreover, this liability level is unique since $\frac{\partial u_i^1}{\partial s \partial L} > 0$.

Step 3: *If $e > \hat{e}$ liability is set to zero $L = 0$ and the planner commits to be more lenient $S_{pp} < S^*$ and $z_{pp} < z^*$.*

If $e > \hat{e}$ from Step 2, we know that $L = 0$. Moreover, if the planner commits to S^* ex ante and to z^* ex post, then the firm has insufficient research incentives $s_{pp} > s^*$. By committing to be more lenient, the planner incurs a second order loss in the adoption and withdrawal standards but this loss is more than compensated by the first order gains generated by the reduction in the rejection standard by the firm.

Proof of Proposition 7

Step 1: *If $F = \infty$ there exists $\hat{e} \in (0,1)$ such that the planner approves at $S = S^*$ and withdraws at $z = z^*$, while the firm replies with $s_i = s^*$ and $S_i = S^*$.*

If $F = \infty$ then for the firm is clearly optimal to acquire information up to the required approval $S_i = S^*$. In fact, the firm controls only s. If $e = 0$ for any S the best reply of the firm is to the left of that of the planner, therefore $s_i < s^*$. If instead $e \to 1$, the value of ex ante experimentation goes to zero, implying that $s_i \to S_i = S^* > s^*$. By continuity there exists $\hat{e}$ such that $s_i = s^*$; moreover, $\hat{e}$ is also unique since $\frac{\partial u_i^1}{\partial s \partial e} > 0$.

Step 2: *If $e \leq \hat{e}$ the planner achieves the first best by requiring $S = S^*$ and committing to $z = z^*$. The firm lies in equilibrium.*

The upper and lower best replies of the firm under costly lying are implicitly defined respectively by

$$\frac{\partial u_i^1}{\partial S_i} - f e^{-S_i} F P(S - S_i) + e^{-S_i} F P'(S - S_i) = 0 \quad (4)$$

$$\frac{\partial u_i^1}{\partial s_i} - a e^{-S_i} F P(S - S_i) = 0. \quad (5)$$

If $e = 0$, the planner can perfectly align the incentives in the bad state by setting $F = \bar{L}$ and requiring $S = 1$ (i.e. $\mathsf{S} = \infty$) which in turns pushes $P(S - S_i) = 1$ and $P'(S - S_i) = 0$. In this case the firm replies by selecting $S_i = S^*$ and $s_i = s^*$. The expected penalty is $FP(S - S_i) = \bar{L}$. As e increases, the planner needs to reduce through S and F the expected penalty since $\frac{\partial u_i^1}{\partial S_i \partial e} > 0$ and $\frac{\partial u_i^1}{\partial s_i \partial e} > 0$. For $e \in (0, \hat{e})$ if the planner keeps $S = 1$ the resulting regulatory regime would be equivalent to the planner committing to L and z under public information. In this case, the firm would not be able to significantly influence the size of the lie since the probability of reaching the report $S = 1$ is zero which implies that reporting $S_i = 1$ is too costly. Thus, simply reducing the fine F will not allow the planner to achieve the first best.

By setting the required report at an interior level, $S < 1$, the planner can effectively control the firm incentives to lie (i.e., $P'(S - S_i) > 0$). For a given S the upper (resp. lower) best reply shift up (resp. left) as F increases. The same comparative statics holds for S, given F. The optimal combination (F,S) solves simultaneously the two equations above evaluated at $s_i = s^*$ and $S_i = S^*$. The firm lies in equilibrium $S_i < S$ because S is interior and $P'(0) = 0$.

Step 3: *If $e > \hat{e}$ then $F = \infty$, there is no lying in equilibrium $S = S_i < S^*$ and the ex post commitment is lenient.*

This follows directly from Proposition 6. For $e > \hat{e}$ the planner shuts liability and commits to a more lenient regulation in both S and z. Under costly lying the planner puts $F = \infty$, thus inducing the informer not to lie $S = S_i$. The expected penalty for lying is then zero, so that the planner controls directly the approval in equilibrium while the firm best replies with s.

A.1 Robustness of Optimal Mix of Instruments under Costly Lying

The planner commits at $t = 0$ to a combination of instruments (S, z, L). Information collection in the first phase is private and non verifiable, so that the firm can at any stage report S and obtain approval. We assume that F, the fine imposed if the firm is found lying, is given and is not an instrument strategically chosen by the planner. Proposition 8 shows the conditions under which the results of Proposition 6 are preserved under private information collection with costly lying.

Proposition 8 **(a)** *With private information collection in the ex ante phase, there exists a level of fine $\tilde{F}$, such that for $F > \tilde{F}$:*

- **(i)** *There exists $\tilde{e} < \hat{e}$ such that the first best is achieved if and only if $e < \tilde{e}$ and all three instruments are used;*
- **(ii)** *Liability is set to zero if $e \geq \tilde{e}$ and is lower than with public information for all e;*
- **(iii)** *In an optimal mix of instruments, the* ***ex post commitment is chosen equal or more lenient*** *than the socially optimal levels z^* and the belief of the firm at approval S_i* ***is lower or equal*** *than the socially optimal approval standard S^*.*

(b) *If $e > 0$, the firm lies in equilibrium: $S_i < S$.*

For a given fine F, changing the standard for approval S has two effects. First it affects the level S_i at which the firm requests approval. Indeed, S_i is increasing in S since the firm reacts to the higher standard by searching more to decrease the size of the equilibrium lie. Second, S also affects the rejection standard s through its effect on the size of the lie. The expected penalty, that is incurred only in the bad state, plays a similar role as a liability rate, by decreasing the expected payoff in the bad state. This explains result (a-ii) that states, as visible in Figure 13 that the liability rate will be smaller than the one imposed in the case with public information, since part of the benefit of reducing excessive research by the firm is provided by the expected fine for lying.

The main result of Proposition 6, that highlights that ex post and ex ante commitments are always lenient in an optimal mix of the three instruments, is preserved under private information

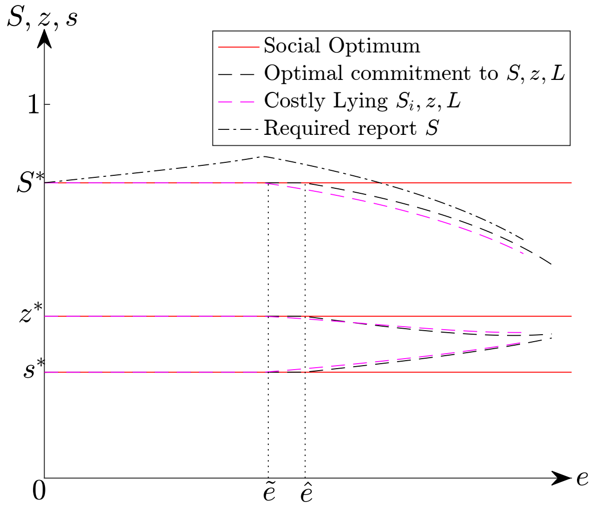

Figure 12: Standards under costly lying.

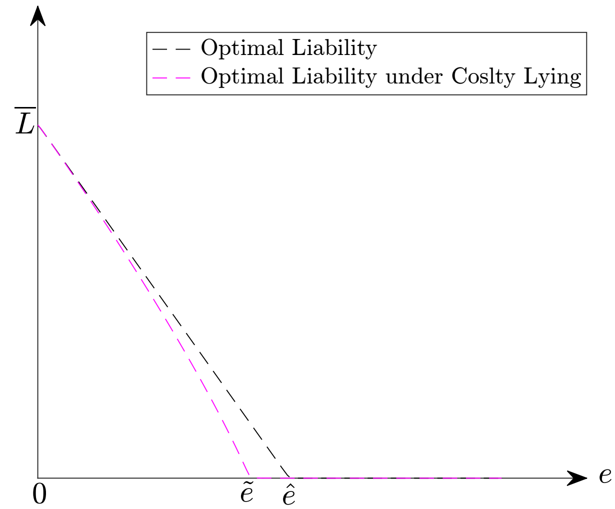

Figure 13: Liability levels under costly lying.

collection when the fine is sufficiently high F. As showed in Figure 12 there is a first phase (for $e < \tilde{e}$) where withdrawal is set at the socially optimal level z^*, the required report standard S is fixed in such a way that the firm asks for approval at the socially optimal level S^* and the liability is chosen in such a way as to prevent socially excessive research and induce a rejection standard at the socially optimal level s^*. At some cutoff value $\tilde{e}$, the first best is no longer attainable and the planner is forced to stop using liability and be more lenient on the approval cutoffs compared to the socially optimal level in order to encourage research at the bottom. Proposition 8 shows that the cutoff externality $\tilde{e}$ is smaller than the cutoff under public information $\hat{e}$, introduced in Proposition 6. Indeed, for $e = \hat{e}$, where under private information the first best is attained with liability set at zero, this is no longer the case under private information since the payoff in the bad state is lower due to the expected penalty for lying, leading to insufficient research at the bottom. The planner is thus forced to decrease the ex post standards to encourage research ex ante.

What does this imply in terms of the optimal level of fine F? In the limit, when $F \to \infty$, there is no lying in equilibrium; welfare converges to the level in the case of public information. In the other extreme, when $F = 0$, the firm can provide any report at no cost and is thus free to introduce the product in the market at any time. This second case is equivalent to the case with public information when the planner has access to only two instruments, z and L, but cannot use ex ante regulation. In between these two extremes, welfare is increasing in F, since as Proposition 8 demonstrates, the fine is a substitute instrument to liability to limit research at the bottom, but the expected fine cannot be set to zero since it also serves to prevent lying. It is thus more efficient to only use liability and set the fine very high in order to prevent misreporting in equilibrium.

Result (b) indicates that the firm always lies in equilibrium. This is directly implied by the fact that the marginal cost of lying at zero is zero, $P'(0) = 0$. Furthermore, in equilibrium the planner knows at what level S_i the firm stopped experimenting, for any standard S required for approval, and thus knows the size of the lie.

Proof of Proposition 8

Step 1: *For any F there exists $\tilde{e}$ such that the first best is achieved if and only if $e < \tilde{e}$.*

If $F = \infty$ then $S_i = S$ there is no lying in equilibrium; the problem reduces to find the optimal mix of instruments when information is public. Thus $\tilde{e} = \hat{e}$ and the result in Proposition 6 hold.

By slightly decreasing F, the firm incentives to lie increase and in equilibrium the firm always lies. At $e = \hat{e}$ we have $S_i < S^*$. In fact, even if $L = 0$, the firm still faces the risk of incurring the penalty for lying; thus, the planner cannot completely shut down all the liabilities in the bad state and achieve the first best as under public information and needs to encourage experimentation. Since the optimal liability when information is public and planner controls all three instruments (i.e. $F = \infty$) decreases in e, there must exist a $\tilde{e} < \hat{e}$ such that the expected fine equals the liability level $L(\tilde{e})$ that the planner would impose under public information. The first best is thus obtained by setting $L = 0$. For all $e \in (0, \tilde{e})$ the planner can reduce liability accordingly in order to align the firm payoff in the bad state under private information with the one under public information.

If $F = 0$ then the firm fully controls approval and the problem reduces to the planner optimally mixing (L, z). Clearly, in this case $\tilde{e} = 0$.

Step 2: *There exists $\tilde{F}$ such that for $F > \tilde{F}$ we have $S_i < S^*$ and $z < z^*$ whenever $e > \tilde{e}$.*

For any F we have that $S_i < S^*$ for $e > \hat{e}$. To see this notice that when $F = \infty$ this holds from Proposition 6. Decreasing F pushes the firm to lie, thus further decreasing S_i.

If $F = 0$ for low externality e the planner might need to be tougher on withdrawal $z > z^*$ to delay approval by the firm. If $F = \infty$ instead the commitment is always lenient $z < z^*$. As F increases ($\tilde{e}$ increases as well). For given $e \in (\tilde{e}, 1)$ by continuity there exists $\tilde{F}$ such that commitment to z is lenient for all $e > \tilde{e}$.

B Appendix B: Planner Benchmark with Costly Withdrawal

In the ex post phase of our baseline model the social planner is able to withdraw (W) and reverse adoption at no cost. In many instances, however, unscrambling the eggs is costly. This appendix extends the model by requiring the planner to pay cost equal to K to exercise the withdrawal option. This extension makes our model applicable to settings such as mergers between companies where reversing the initial decision is costly. By introducing a fixed reversibility cost we can also connect our Wald problem with reversible decision (where $K = 0$) to the classic Wald problem with irreversible decision ($K = \infty$).

In the ex post phase, the planner solves the optimal stopping problem

$$\hat{u}_p^2(\sigma_0) \equiv \sup_{\tau} \mathbb{E}_{\Omega\times\Theta}\left[-e^{-r\tau}K + q_\tau\left(\int_0^\tau e^{-rt}v_p^G dt\right) + (1-q_\tau)\left(\int_0^\tau e^{-rt}v_p^B dt\right)\middle|\, \sigma_0\right]$$

for $\theta \in \{G, B\}$ with $K \geq 0$, where $\tau^* = \infty$ if $K \geq -v_p^B/r$.

Proposition 9 **(a)** *The withdrawal standard $z^*(K)$ is decreasing in the fixed cost, K.*

(b) *The rejection standard $s^*(K)$ and the adoption standard $S^*(K)$ are both increasing in K.*

(c) *If $K \geq -v_p^B/r$ all the optimal standards coincide with the solution to the classic Wald problem.*

The intuition for (a) is straightforward. According to part (b), the ex ante standards increase in K. The higher the withdrawal cost, the less valuable the withdrawal option for the planner, so that the planner becomes more careful before adopting (S^* increases) and rejects more frequently (s^* increases). Figure 15 shows the value function at the optimal three-thresholds policy (s^*, S^*, z^*). Notice that Proposition 1-(c)'s property (that the ex ante optimal standards are centered around the withdrawal threshold z^*) does not necessarily hold when $K > 0$.

According to part (c), when the cost of withdrawal is sufficiently high, we have $z^* = 0$ and the planner never reverses adoption, as in the classic Wald problem with irreversible adoption. Figure 14 displays the classic Wald value function in red and the ex post value in blue which is then in linear in q. As K is reduced below $-v_p^B/r$, the planner withdraws at positive beliefs $z^*(K) > 0$; the ex post value becomes convex in q, as shown in Figure 15. In the limit when withdrawal is costless ($K = 0$) we are back to our baseline model, as shown in Figure 1.

B.1 Best Replies Construction

As also explained in the main text, the optimal solution of our two-phase experimentation problem can be represented in terms of thresholds. The solution thresholds are the belief levels at which the smooth-pasting conditions in all experimentation phases are satisfied. The optimal adoption standard S, for instance, solves the corresponding optimality condition for a given rejection s and optimal withdrawal z^* standards. Solving simultaneously also for the optimal

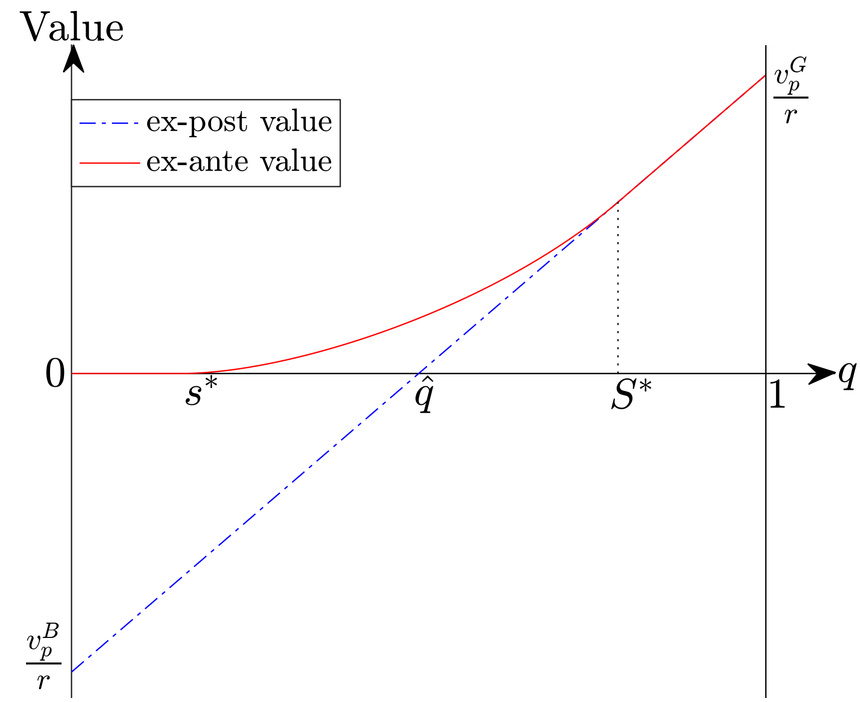

Figure 14: Value function for classic Wald problem.

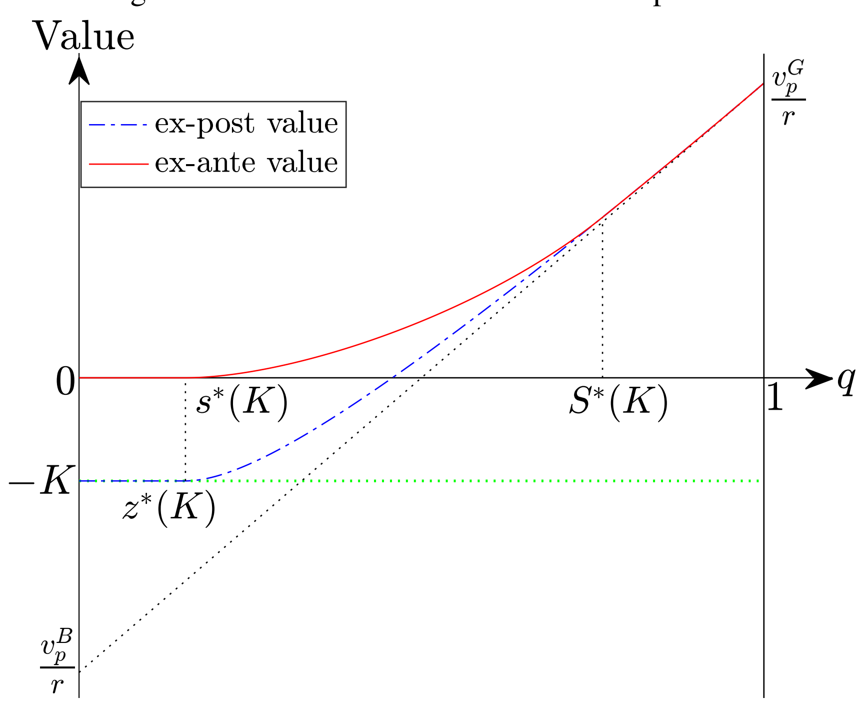

Figure 15: Value function with withdrawal cost K.

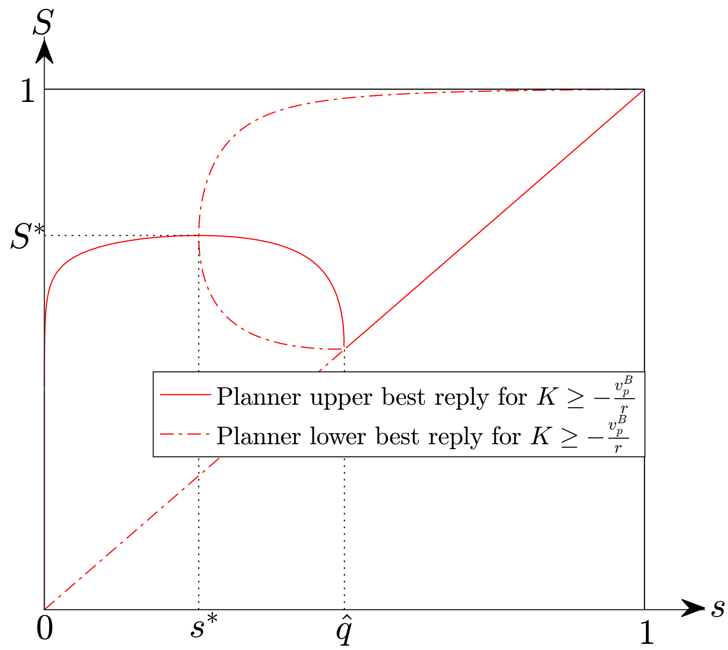

Figure 16: Best replies for classic Wald problem.

rejection standard s^* given S and z^*, the optimal solution is found at the intersection of the upper best reply $S = B(s)$ and the lower best reply $s = b(S)$ given the optimal withdrawal z^*.

To connect the optimality conditions in the classic Wald problem (with $K = \infty$) to the solution of our model with reversible adoption (with $K = 0$), we now analyze the impact of the cost K of exercising the withdrawal option on these best replies. When the withdrawal cost K is sufficiently high, the upper and lower best replies are as in Figure 16 and the optimal policy (s^*, S^*) at their intersection is the well-known solution of the classic Wald problem.[43] Given that withdrawing is too costly, in this case the planner never withdraws (i.e. $z^* = 0$). Figure 14 shows the resulting value ex ante function along with the ex post value function which is linear in q in this specific case.

When K is sufficiently low, the planner optimally withdraws as soon as the belief hits the withdrawal standard $z^*(K) > 0$. The continuation value becomes convex in q, as displayed in Figure 15. The fact that, differently from the classic Wald, the continuation value upon adoption is not linear influences the shape of the upper best reply $B(s)$. To understand this new shape, we need to analyze the continuation value upon adoption.

Proposition 10 *Given any commitment to a suboptimal withdrawal standard $z \neq z^*$ there exists a fixed withdrawal cost $K > 0$ such that when optimizing (taking K into account) the continuation value upon approval is the same as the value under commitment.*

Figure 17 highlights the result in Proposition 10. The dashed-dotted line tangent at z^* to the

[43] See Henry and Ottaviani's (2017) Appendix B.

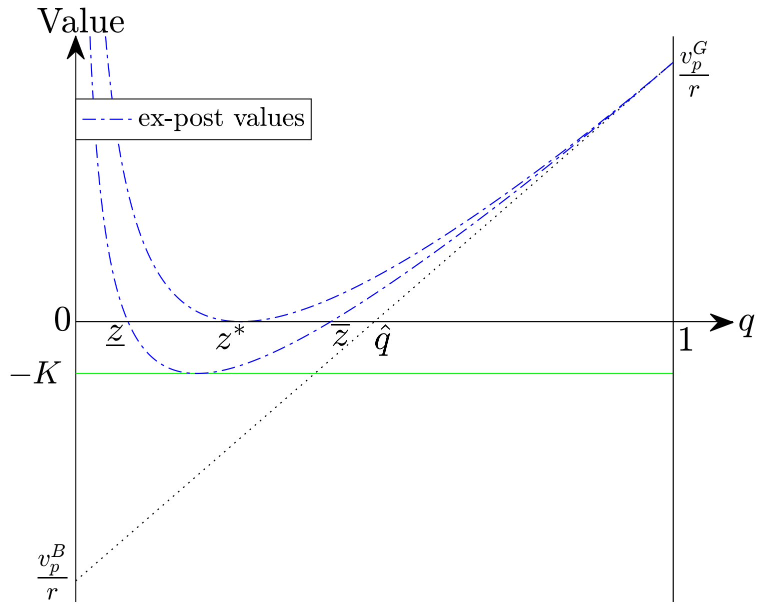

Figure 17: Value under commitment or with costly withdrawal.

horizontal axis represents the optimal continuation value when $K = 0$. If the planner commits to a suboptimal withdrawal such as $\underline{z} < z^*$ or $\overline{z} > z^*$, the continuation value shifts down. Clearly, there exists a withdrawal cost $K > 0$ such that when solving for the optimal withdrawal $z^*(K)$ the continuation value coincides with the value resulting under commitment and $K = 0$.

As described in Proposition 1, the optimal solutions to this problem consist in a triplet of standards (s^*, S^*, z^*), each one solving the corresponding smooth-pasting condition given the other standards. To understanding the shape of the upper best reply in Figure 19, Figure 18 shows an example of this solution. Fix $s_1 < s^*$and $z^*(K)$, the smooth-pasting condition for the adoption standard admits two solutions. In particular, there are two value functions tangent to the continuation value respectively at $\underline{S_1}$ (from below) and at $\overline{S_1}$ (from above), both solving the smooth-pasting condition for a given s_1. To find the adoption standard that best responds to s_1, we keep $\overline{S_1}$ (a point of maximum) but exclude $\underline{S_1}$ (a point of minimum).[44] Thus, $\overline{S_1} = B(s_1)$ defines the upper best reply to s_1.

Repeating the process for any s given $z^*(K)$ we plot the upper best reply function $S = B(s)$ in Figure 19, along with $B(s)$ of the classic Wald problem. When the withdrawal cost is not excessive, the upper best reply shifts down and becomes loop-shaped around $\underline{z}$ and $\overline{z}$. Fixing s_1 on the horizontal axis, Figure 19 shows the two solutions $\underline{S_1}$ and $\overline{S_1}$ to the optimality condition for the adoption standard; these solutions are generated by the loop-shaped part of the upper best reply, as in Figure 18.

As K decreases the loop shrinks and closes up at z^* when $K = 0$, as illustrated in Figure 20.

[44]The second-order conditions given by Arkin and Slastnikov (2013) allow us to discard solution $\underline{S_1}$ as a point of minimum.

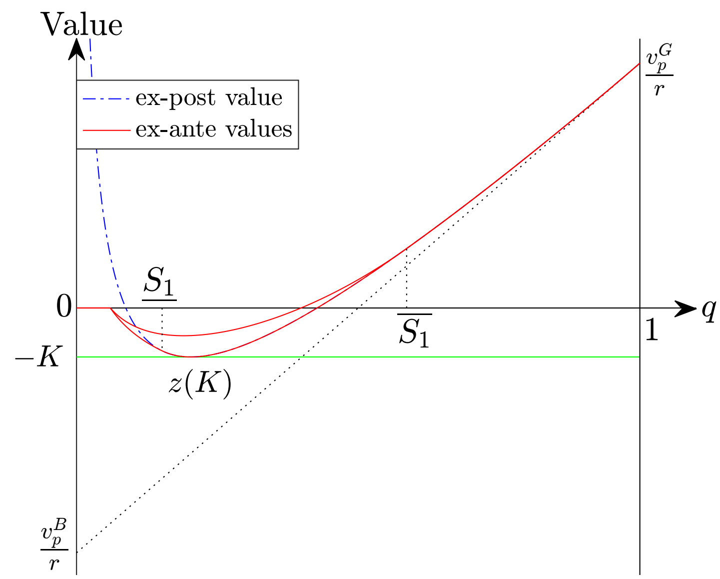

Figure 18: Value at multiple solutions.

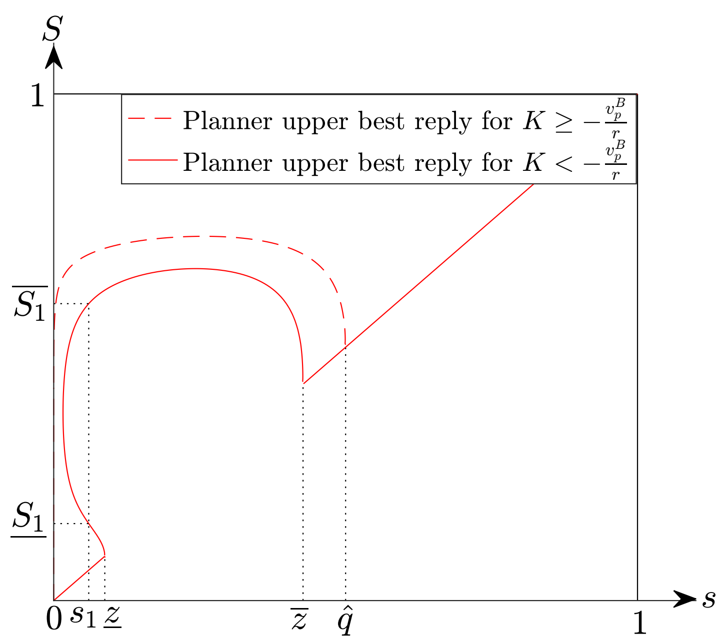

Figure 19: Upper best replies at $F \geq \frac{v_p^B}{r}$ and $F < \frac{v_p^B}{r}$.

The highest dotted curve is the upper best reply for high withdrawal cost K corresponding to the classic Wald with irreversible decision, when $z^* = \underline{z} = 0$ and $\overline{z} = \hat{q}$. As K decreases, the upper best reply becomes loop-shaped and the planner withdraws at the positive belief $z^*(K) \in (\underline{z}, \overline{z})$. At $K = 0$, as in our baseline model, the loop closes up at the optimal withdrawal level z^*. Finally, the green line represents the path of the optimal ex ante rejection and adoption standards as a function of K. Consistent with Proposition 9-(b), the path $(s^*(K), S^*(K))$ is increasing in K.

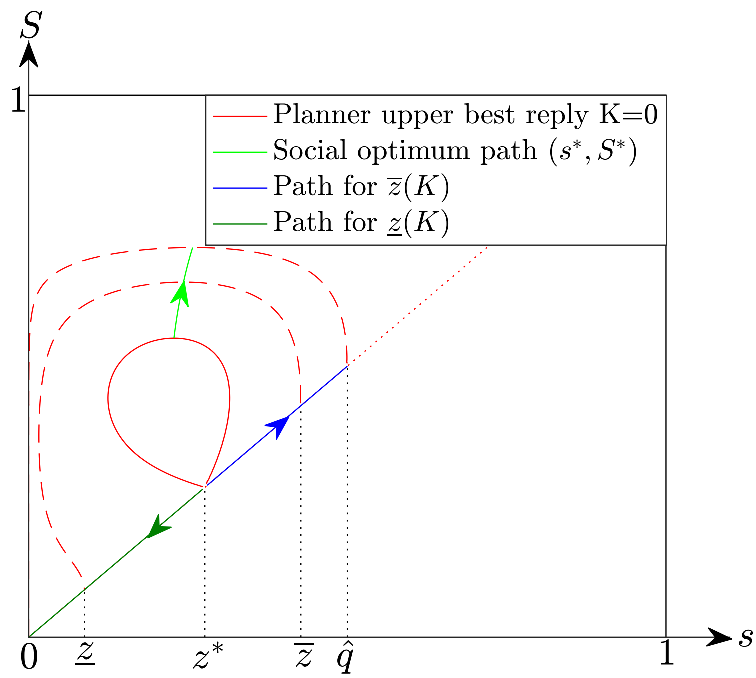

Figure 20: Ex ante standards path as function of K.

In our baseline model we have $K = 0$, thus when plotting the upper best reply in Figure 2 we need to exclude the lower part of the loop, corresponding to the solutions to the smooth pasting that are points of minimum.

B.2 Proofs for Appendix B

Proof of Proposition 9

Step 1: *The optimal withdrawal standard z^* is decreasing in K.*
If the adoption decision has been taken at belief S, we can write the planner problem as

$$\max_z \left\{ \frac{e^S}{1+e^S} \left[\left(\frac{v_p^G}{r} \right) - \left(\frac{v_p^G}{r} + K \right) \psi(S,G,z) \right] + \frac{1}{1+e^S} \left[\left(\frac{v_p^B}{r} \right) - \left(\frac{v_p^B}{r} + K \right) \psi(S,B,z) \right] \right\}.$$

Taking first-order condition and rearranging we have $z^*(K) = -\log \frac{r_2(v^G+rK)}{r_1(v_p^B+rK)}$ from which $\frac{\partial z^*(K)}{\partial K} = \frac{r(v_p^G - v_p^B)}{(v_p^G+rK)(v_p^B+rK)} < 0$ for $K < -\frac{v_p^B}{r}$ which is the relevant region for interior solutions.

Step 2: *The optimal ex ante standards s^* and S^* are increasing in K.*
The first order condition with respect to S^* is

$$u^1_{p,S}(s,S,\sigma) = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \Psi \left\{ \begin{array}{c} f\left[u_p^2(S,G,z) + e^{-S} u_p^2(S,B,z) + (1+e^{-S})\frac{c}{r}\right] \\ +g(1+e^{-s})\frac{c}{r} - e^{-S}\left[\frac{c}{r} + u_p^2(S,B,z)\right] \end{array} \right\},$$

where $u_p^2(S,G,z) = \left(\frac{v_p^G}{r}\right) - \left(\frac{v_p^G}{r} + K\right) e^{-r_2(S-z)}$, $u_p^2(S,B,z) = \left(\frac{v_p^B}{r}\right) - \left(\frac{v_p^B}{r} + K\right) e^{r_1(S-z)}$ and $f < 0$ and $g > 0$.[45] By the implicit function theorem we have that

$$\frac{\partial S^*(K)}{\partial K} = -\left.\frac{\frac{\partial u_{p,S}^1}{\partial K}}{\frac{\partial u_{p,S}^1}{\partial K}}\right|_{S=S^*}.$$

The denominator is negative being S^* a point of maximum, hence we can focus on the sign of the numerator. After some computations we have

$$\frac{\partial u_{p,S}^1}{\partial K} = \frac{e^{\sigma_0}}{1+e^{\sigma_0}} \Psi \left\{ f \left[u_{p,K}^2(S,G,z) + e^{-S} u_{p,K}^2(S,B,z) \right] - e^{-S} u_{p,K}^2(S,B,z) \right\},$$

where the term

$$u_{p,K}^2(S,G,z) + e^{-S_p} u_{p,K}^2(S,B,z) = \begin{pmatrix} -e^{-r_2(S-z)} - e^{-S} e^{r_1(S-z)} \\ + \left[u_{p,z}^2(S,G,z) + e^{-S} u_{p,z}^2(S,B,z) \right] \frac{\partial z(K)}{\partial K} \end{pmatrix}$$

is negative because by the Envelope Theorem we have $u_{p,z}^2(S,G,z) + e^{-S} u_{p,z}^2(S,B,z) = 0$ at $z = z^*$. Moreover, $u_{p,K}^2(S_p,B,z_p) = -e^{r_1(S_p - z)} + r_1 \left(\frac{v_p^B}{r} + K\right) e^{r_1(S_p - z_p)} \frac{\partial z(K)}{\partial K} < 0$ at any interior solution z^*. Overall, $\frac{\partial u_{p,S}^1}{\partial K} > 0$ proving that S^* is increasing in K. Using the same argument one can show that also $s^*(K)$ is increasing in K.

Step 3: *If $K \geq -\frac{v_p^B}{r}$ then the $z^* = 0$ the ex ante standards are the one solving the classic Wald problem.*

In this case the cost of exercising the withdrawal option is higher than the whole flow of profit resulting in the bad state of the world. Clearly, for the planner it is never optimal to select an interior withdrawal threshold. Therefore, the planner never withdraws, $z^* = 0$. The ex post value at $z^* = 0$ (corresponding to $\mathrm{z}^* = -\infty$) is

$$u_p^2(S,0) = \frac{e^{\mathrm{S}}}{1+e^{\mathrm{S}}} \left(\frac{v_p^G}{r}\right) + \frac{1}{1+e^{\mathrm{S}}} \left(\frac{v_p^B}{r}\right),$$

a straight line in the regular belief space q. The planner ex ante solves a classic Wald problem for which abandoning results in a payoff of 0 and irreversibly adopting results in an expected payoff of $u_p^2(S,0)$.

Proof of Proposition 10

Assume $K = 0$, then the continuation value upon adoption at $S = \sigma$ is

$$\hat{u}_p^2(\sigma) = \frac{e^{\sigma}}{1+e^{\sigma}} \left(\frac{v_p^G}{r}\right) \left(1 - e^{-r_2(\sigma - \mathrm{z}^*)}\right) + \frac{1}{1+e^{\sigma}} \left(\frac{v_p^B}{r}\right) \left(1 - e^{r_1(\sigma - \mathrm{z}^*)}\right).$$

[45] Henry and Ottaviani (2017) characterize the properties of the functions Ψ, ψ, f, and g.

Notice that $\hat{u}_p^2(\sigma) \geq 0$ for any σ, decreasing for $\sigma < z^*$ and increasing for $\sigma > z^*$. At $\sigma = z^*$ the continuation value is tangent to the horizontal axis and $\hat{u}_p^2(z^*) = 0$. If our social planner commits to a suboptimal $z \neq z^*$ then $\hat{u}_p^2(\sigma)$ shifts down in the belief space (see the blue dotted lines in Figure 17). Therefore, there are two beliefs $\underline{z}$ and $\bar{z}$ such that $\hat{u}_p^2(\mathrm{z}) = 0$ for $z \in \{\underline{z}, \bar{z}\}$ where either one of those is the suboptimal commitment and $z^* \in (\underline{z}, \bar{z})$.

If instead $K > 0$ then the continuation value upon adoption at $S = \sigma$ is

$$\hat{u}_p^2(\sigma) = \frac{e^\sigma}{1+e^\sigma}\left(\frac{v_p^G}{r}\right)\left(1 - e^{-r_2(\sigma - \mathbf{z}^*(K))}\right) + \frac{1}{1+e^\sigma}\left(\frac{v_p^B}{r}\right)\left(1 - e^{r_1(\sigma - \mathbf{z}^*(K))}\right) \tag{6}$$

$$-K\left(\frac{e^\sigma}{1+e^\sigma}e^{-r_2(\sigma - \mathbf{z}^*(K))} + \frac{1}{1+e^\sigma}e^{r_1(\sigma - \mathbf{z}^*(K))}\right). \tag{7}$$

Notice that $\hat{u}_p^2(\sigma) \geq -K$ for any σ, decreasing for $\sigma < z^*(K)$, and increasing for $\sigma > z^*(K)$. At $\sigma = z^*(K)$ the continuation value is tangent to the horizontal level $-K$ and $\hat{u}_p^2(z^*(K)) = -K$.

Given $K = 0$ and any suboptimal commitment $z^c \neq z^*$, the withdrawal option cost $K > 0$ that solves

$$l(\sigma, K) \equiv \hat{u}_p^2(\sigma, z^*(K)) - u_p^2(\sigma, z^c) = 0$$

is such that the optimal solution taking into account the continuation value is the same as the solution under commitment z^c and no cost, $K = 0$. Notice that $\frac{\partial l(\sigma,K)}{\partial K} < 0$, $l(\sigma, 0) > 0$, and $l(\sigma, K) \leq 0$ for K sufficiently high for any σ. Hence, given z^c there exists $K > 0$ such that $l(\sigma, K) = 0$.