Strategic Experiments Under Regulatory Uncertainty*

Jenny S. Kim†

Abstract

I present a model of policymaking in complex domains. I apply the model to a hypothetical situation in which a firm's product has uncertain social impacts. The firm can acquire more information about this, but knows that this information will be "public" in the sense that it will also be observed by the regulator. The firm's choice about information is represented as a Blackwell experiment. After the result of the experiment is realized, the firm and regulator can each take a unilateral costly action to reveal the information with certainty. I show that the firm-optimal can be achieved by providing a binary experiment that is informative only about whether the firm or the regulator has an incentive to take a costly action. I further extend the model to allow the firm and regulator to bargain prior to the game and allow the firm to have private information about the product.

*I thank John Patty, Maggie Penn, Jessica Sun, Kyungmin (Teddy) Kim, Carlo Prato, Ryan Hubert, Barton Lee, Peter Buisseret, Lauren Futter, Allison Stashko, Luis Martinez, Alex Hirsch, Brenton Kenkel, Mike Gibilisco, and Sangmin Lee, as well as the audience members at APSA 2023 Annual Meetings in LA, California, Aspen 2024 Conference on Political Economy, Theory & Models Working Group at Emory University, 2024 Formal Theory Conference at Washington University in St. Louis, and APSA 2024 Annual Meetings in Philadelphia, for very helpful comments on this project. All errors are my own.

†Departments of Political Science, Emory University. Email: *seoyeon.kim@emory.edu*.

1 Introduction

Emerging technologies continuously reshape markets, often adding new regulatory concerns to an industry. In response, regulatory agencies are tasked with evaluating these new developments with regard to the existing regulatory practices and determining potential violations. Yet the significant complexity of many industries, coupled with the inadvertent co-production of latent hazards by independent firms, leaves both firms and agencies uncertain about the new product's expected safety or, more broadly, its impact on society's welfare (Carpenter, 2004; Carpenter and Ting, 2007; McCarty, 2017; Montagnes and Wolton, 2017; Sabel, Herrigel and Kristensen, 2018).

For instance, part of the reason behind the grounding of the Boeing 787 Dreamliner fleet was the fact that defects that could emerge in battery manufacturing were not caught by inspection (NTSB, 2013); the regulation of a decongestant drug pseudoephedrine highlights how a product can become a critical component in harmful activities, such as the production of meth, illustrating how regulatory concerns can emerge in unforeseen ways. In other cases, a lot is still unknown about whether the facial recognition technology (FRT)'s benefits will outweigh the potential problems regarding misidentification, underscoring the challenges agencies face in anticipating and addressing latent risks within existing regulatory frameworks. With the U.S. Supreme Court's recent ruling in *Loper Bright Enterprises v. Raimondo* endowing the courts with the power to rely on their own interpretation of ambiguous laws, the legal landscape has recently become far more uncertain for both firms and the agencies that regulate them.

I consider the firm's optimal information acquisition in this context. In the theoretical framework, both the firm and the regulator are uncertain about the true social impact of the firm's product. Uncertainty over the outcome can be fully resolved through players' decision to take the issue to court, but prior to the decision, the firm conducts a one-time observable Blackwell experiment that is informative about the truth. After the firm and regulator observe the experiment's outcome, each chooses whether to accept the status quo level of regulation or pay a cost to go to court and enforce a new regulatory standard that reflects the "true" social impact of the product.

It is straightforward to see that players will go to court only if they expect to gain by doing so. The firm thus wants to learn more about the product, as this helps the firm make an informed decision about whether or not to risk going to court. Importantly, however, the firm's choice of an experiment is *public* in my model. For example, in the real world, how tech companies such as Apple, Google, Meta, and X/Twitter deploy new features using AI shapes not only their own beliefs about the products' social, economic, and political impacts - it also shapes the beliefs held by outsiders, such as regulators and politicians; oil companies shape their own and the public's beliefs through their choice of which experts to hire to evaluate new oil spill prevention technologies; pharmaceutical companies choose which research design to use when seeking approval of their drugs, and so forth. In the context of the model, this implies that both the firm and the regulator update their beliefs about the product's social impact via the experiment. In other words, the regulator—whose preferences are not fully aligned with the firm—may also benefit from the firm's information acquisition. This paper focuses on how such "learning in front of an audience" shapes the firm's incentives to provide information and fully characterizes the firm's optimal signal structure.

I show that the firm's optimal experimental design in this framework manipulates the observability of the underlying state to the firm's advantage. Despite having complete flexibility in designing its information acquisition process, the firm's optimal choice is a simple binary experiment: the outcome of the experiment effectively signals only whether the regulator or the firm may have an incentive to go to court. More formally, there are essentially two classes of binary experiments that can be optimal for the firm that reveal whether the product quality is *above* or *below* some threshold.

Specifically, I characterize the optimal experiment by examining three different sub-cases. When the quality of the product is likely to be high, i.e., the firm's risk of going to court is low, the firm can perfectly exploit its information-designing power by revealing all the good information about the product and concealing the rest. This is equivalent to commissioning a study that reveals only whether the product is "extremely safe" or "not extremely safe." Here the prior belief about the product is favorable enough that the firm can perfectly separate out the states under which it would challenge the status quo ("extremely safe"); and the key is that even when the study finds that the product is "not extremely safe," the regulator still believes that the product must be safe *enough* and refrains from going to court.

With moderate risk, the firm-optimal still involves only revealing good information about the product, but now the firm is unable to reveal all of them as in the low-risk case. This is because the regulator's prior belief about the product is moderate that if the study reveals that the product is not "extremely safe," the regulator is now convinced that it must be unsafe and goes to court. For the firm to deter the regulator from going to court after a bad signal, it optimally has to pool some of the good information with the bad information. By doing so, the firm induces the regulator's posterior belief after the "not extremely safe" outcome to be just high enough that he is persuaded not to take the issue to court.

When the risk of going to court is high, the firm optimally reveals whether the product is extremely harmful to the society, i.e., commissions a study that reveals whether the product is "extremely *un*safe" or "not extremely *un*safe." Here the regulator is very strongly incentivized to go to court in the first place. In order for the firm to dissuade the regulator from going to court, the firm has to commit to disclosing some of its worst states ("extremely unsafe"), after which the regulator goes to court and tightens the regulation. The firm ex-ante gains by doing so because this way, the regulator's posterior belief conditional on the product being "not extremely unsafe" is just high enough to deter him from going to court. This last result can in particular be helpful in explaining why we often see firms voluntarily disclosing damaging information about its own product. For example, Nike has implemented a Responsible Disclosure Program that encourages external researchers to report security vulnerabilities in its products; Meta has hired independent audit boards that are tasked with revealing controversial practices related to misinformation, harmful content, and privacy issues. We expect to observe these behaviors when the firm is *least* confident about the potential social impact of the product.

The model further shows that both players sometimes have shared preferences for silence. In equilibrium, the firm strategically manipulates information to maximize the regulator's silence; in other words, by fully exploiting its informational advantage, the regulator is persuaded not to pursue court action even when it could ex-post successfully challenge the status quo. As a result, the regulator often ends up worse off due to the firm's information generation. I find that the regulator sometimes prefers to let this happen. The idea is that as the firm gains more information about the product, it becomes better at identifying *when* it is strategically advantageous to go to court and in turn goes to court less often compared to when it lacks information. The regulator, in turn, may benefit from the firm's increased reluctance to go to court and willingly forgoes its own opportunities to pursue legal action in favor of maintaining silence.

This framework is related to an extensive and growing literature on regulation in incomplete information settings. Previous works consider an optimal stopping problem in the context of pre-market approval where an agency learns about the policy environment through the firm's experimentation (e.g., Carpenter, 2002, 2004; Carpenter and Ting, 2007; Carpenter, Grimmer and Lomazoff, 2010; Henry and Ottaviani, 2019; McClellan, 2022) or their efforts at self-regulation (McCarty, 2017). Other works examine how market competition interacts with politics in innovative markets (Callander, Foarta and Sugaya, 2022*a*,*b*), or emphasize the bureaucratic problems inherent to delegation when new regulatory concerns emerge (Boehmke et al., 2006; Gailmard and Patty, 2012; Montagnes and Wolton, 2017).

Much of the work to date has understandably considered settings in which players' payoffs are conditional on each other's action: the firm decides how much experimentation to perform on a product; after observing the outcome, the firm decides whether to submit the product for review. The regulator observes the information, updates his beliefs, and then decides whether or not to approve the product (Carpenter, 2002, 2004; Carpenter, Grimmer and Lomazoff, 2010). In this setting, the firm must apply for the product in order for the regulator to review it. This means that information can never backfire for firm. If the firm's experiment returns a bad outcome, and if the firm decides not to submit the product, the regulator is unable to reject something that has not been submitted. The firm's primary incentive here is thus to use information to persuade the regulator to approve the product whenever the firm submits it for review.

The key difference of my model is the firm and the regulator's ability to act *independently* after acquiring information. We often observe examples of this within the context of post-market complexities associated with a firm's product. In these scenarios, the regulator can take unilateral measures that tighten the preexisting regulation, which implies that now information can *backfire*. If the outcome of the experiment is damaging for firm, the regulator can use that information to shift the regulatory landscape without having to condition its action on firm's submission. Now the firm needs to consider how much to learn about the product in light of the fact that the regulator also observes the information and can use it against the firm.[1]

More broadly, my model contributes to the large body of research that highlights how strategic information provision is affected by the preferences of the decision-maker (e.g. Callander, 2011*a*,*b*; Gailmard and Patty, 2017; Bils, Carroll and Rothenberg, 2020; Libgober, 2020; Prato and Turner, 2022) and, concomitantly, how decisions are informed and chosen when faced with strategically information provision (Austen-Smith, 1993; Potters and Van Winden, 1992; Callander and Hummel, 2014; Callander and Harstad, 2015; Schnakenberg, 2017; Schnakenberg and Turner, 2019, 2021, 2023; Awad, 2020; Ellis and Groll, 2020; Patty and Penn, 2022; Dellis, 2023; Patty and Penn, 2023). For instance, Patty and Penn (2022) find that the employer might optimally withhold sensitive information that can be used for discrimination. In Patty and Penn (2023), the designer might want to incorporate noise into his decision in order to manipulate behavior. These results resonate with the firm in my model committing to receiving a noisy signal.

2 Benchmark Model

Consider two players, Firm (she) and Regulator (he). Firm wants to expand her marketing claims for a product. The product is currently subject to a certain level of regulation, with Firm and Regulator's utility from the regulation reflected in the status quo division $(x, 1-x)$ where $x \in [0,1]$.

There is currently only preliminary evidence about the product's expected social im-

[1] Note that this model can still account for the process of product review or bargaining. So long as both players retain the option to independently challenge the status quo regulation later down the road, the qualitative results of my model remain consistent.

pact, ω. This underlying state conveys the extent to which Regulator can impose rules or restrictions on the product within legal confines. To keep the model as transparent and simple to present as possible, I only focus on the interval over which Firm and Regulator have misaligned interests and assume that Firm always prefers a more lenient regulation than Regulator does. Both players can decide to pay some cost to go to court and have an (unmodeled) court reveal ω and enforce a new regulation based on the underlying state.

Prior to this decision, Firm chooses a publicly observable experiment that generates some information about state ω. This outcome becomes the data for both players' decisions that follow. We could expect two counteracting incentives to prevail: one where Firm wants to learn more about the product so that she could make an informed decision about whether or not to challenge the status quo, and another where she wants Regulator to be sufficiently unaware that he is persuaded not to go court even when the current level of regulation is too lenient relative to the "true" social impact of the product. Generating less information can thus be strategically optimal for Firm. Each stage of the game is detailed below.

Initial Stage. Nature first draws ω from distribution F with expectation $\mathbb{E}[\omega]$. I assume that F has continuous and strictly positive density f on $[0,1]$. State ω is not directly observed by any player. Players then observe the status quo division $(x, 1-x)$ and the cost of going to court $c \in [0,1]$, which is common knowledge to both players.

Signal Stage. Firm engages in a publicly observable Blackwell experiment that informs both players about the outcome. The experiment gathers more information about ω. Firm can choose how *much* to learn about it, and any information that Firm learns is accessible to Regulator. Players condition their strategy in the conflict stage on such endogenously acquired public information. Formally, Firm can choose any ordered set S and joint distribution v over $[0,1] \times S$. Players observe the realization $s \in S$ and Bayesian updates their beliefs. Let F_s denote the distribution that represents the posterior belief following the realization $s \in S$.

Conflict Stage. Players simultaneously decide whether or not to go to court and enforce a new regulation based on the public signal they have received in the signal stage.[2] In this stage, they choose either to go to court (C) or not (N). If both players choose not to go to court and stay silent, they receive the initial status quo payoffs. If either of the players chooses to go to court, the true state is realized and players' realized payoffs are $(\omega, 1-\omega)$. Intuitively, Regulator goes to court if the outlook of going to court post-signal is advantageous enough for himself; Firm may also do so if she is fairly confident about the outcome and believes that she will be able to challenge the status quo. Further, a player choosing to go to court pays a net cost c, which measures the general effort needed for the true state to be realized. The cost implies that players won't be willing to overturn the initial status quo division unless the additional benefit they expect to gain by going to court exceeds its cost. The sequence of the game is as follows:

1. Nature draws a true state $\omega \in [0,1]$ from distribution F.
2. Players observe the cost of going to court c and the initial division of dollar $(x, 1-x)$.
3. Firm chooses a publicly observable experiment defined by set S and distribution v.
4. Players observe the realized signal $s \in S$.
5. Players play a simultaneous game.

Firm \ Regulator	N	C
N	$(x, 1-x)$	$(\omega, 1-\omega-c)$
C	$(\omega-c, 1-\omega)$	$(\omega-c, 1-\omega-c)$

Table 1: Payoff Structure

Below I offer a few comments on this model's assumptions. First, in the model, the incentives of Firm are misaligned with those of Regulator. In particular, I assume that Firm does not suffer the full social damage even when the product is revealed to be socially harmful

[2] Note that the results of the model are identical even if we assume that the players move in sequential order. This is because players never go to court at the same time in equilibrium. In such a setup, one of the players moves first and decides whether to go to court or stay silent; the other player then also chooses between the same actions. I assume throughout for simplicity that players act simultaneously.

and does not necessarily recoup the full social benefits (Henry, Loseto and Ottaviani, 2022). Firm thus simply wants to push the regulatory boundaries and promote her new product under fewer regulations ($\omega = 1$). Regulator, on the other hand, believes that there are potential risks that have not been fully explored and wants to impose stricter standards ($\omega = 0$). Alternatively, he is cautious about setting a strong precedent that may limit his future ability to regulate Firm, or he prioritizes guarding his reputation for protecting consumer safety (e.g. Carpenter, 2004).

Note that Regulator need not have direct preferences against the product, i.e., we should not think of $\omega = 0$ as the state where the product is harmful. Instead, he has preferences for some level of regulation on the product with uncertain risks. In this sense, we may also assume Regulator who simply prefers to impose fewer regulations when the product has high quality and strict regulations otherwise. The model could then be interpreted as a subgame of this game where the quality of the product is "low," and Firm and Regulator are in conflict over the level of regulation. Firm still wants the product to be subject to lenient regulations while Regulator prioritizes safety over profitability and prefers it to be regulated. In Section E.2, I directly explore this with a different payoff matrix that reflects Regulator's preference for accurately matching the regulation with underlying state ω.[3]

Second, both players are symmetrically uninformed about the true social impact of the product. This value is initially unknown to both players and remains so unless either of the players pays some cost and goes to court, reflecting the notion that firms may be uncertain as to the quality and safety of their products (Carpenter, Grimmer and Lomazoff, 2010). This may especially be the case given that players in my model act anticipating the future realization of truth via legal proceedings, which are influenced by various external factors such as changes in the political climate or public perception. Such inherent unpredictability makes it challenging for both Firm and Regulator to forecast the results of a potential dispute. I relax this in Section 6.2 and discuss how the results change when we assume privately informed Firm.

[3]I consider a binary example where $\omega \in \{0, 1\}$ and find that Firm optimally provides no information more often than in the benchmark model. This is coming from the fact that Firm may now *want* Regulator to go to court so that she can free-ride on the cost while changing the status quo to her favor. Otherwise the key qualitative outcomes remain unchanged.

Third, I interpret the cost of going to court c as the "enforcement cost" required to contest the existing regulation and establish a new standard. For simplicity c is assumed to be common for all players; however, note that this enforcement cost may have to be interpreted differently for Firm and Regulator, especially since Firm is likely to have the technology to fully reveal the state and communicate it to others with no costs. We could then alternatively think of heterogeneous costs for Firm and Regulator where $c_R = c_F + \gamma$, and γ is regulator's cost to *learning* the state and c is administrative costs, i.e., Firm has no costs of learning. I characterize this in Appendix F and find that the results are qualitatively similar.

Lastly, the current setup of the model allows going to court as the only action available to the players after the experiment and the detailed process of the legal action is modeled in a "black-box" fashion. While this modeling choice was made for simplicity, the model can easily account for other processes of post-experiment bargaining or product review without changing the core results of the model. In particular, Firm's optimal choice of experiment will be identical so long as the conflict stage exists; this is because Firm's incentive to learn more/less depends primarily on the fact that both Firm and Regulator can unilaterally pay some cost to reveal the underlying state and enforce a new regulation.

3 Equilibrium Analysis

This section provides a full characterization of the model. I first solve the last stage of the game where both Firm and Regulator, with some belief about ω, decide whether or not to pay some cost to go to court and enforce the underlying state. Then, I analyze Firm's optimal learning problem by varying the ex-ante level of product risk.

3.1 Decision in the Conflict Stage

Note that in the conflict stage, players are still uncertain about the true value of ω. Their decision in this stage depends on their belief about ω given realization s from the experiment, which means that we can rewrite the game in this stage as in Figure 1a.

Firm \ Regulator	N	I
N	$(x, 1-x)$	$(\mathbb{E}[\omega\|s], 1-\mathbb{E}[\omega\|s]-c)$
I	$(\mathbb{E}[\omega\|s]-c, 1-\mathbb{E}[\omega\|s])$	$(\mathbb{E}[\omega\|s]-c, 1-\mathbb{E}[\omega\|s]-c)$

(a) Payoff Structure

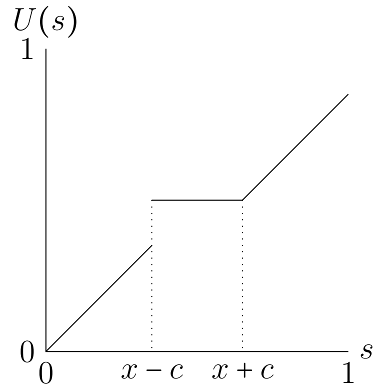

(b) Firm's Indirect Utility

Figure 1: This figure illustrates the payoff structure (left panel) and Firm's indirect utility (right panel) in the conflict stage.

Solving for the Nash equilibrium for the conflict stage,

- Firm remains silent and Regulator goes to court if $\mathbb{E}[\omega|s] < x - c$,
- Firm goes to court and Regulator remains silent if $\mathbb{E}[\omega|s] > x + c$,
- Both remain silent if $x - c \leq \mathbb{E}[\omega|s] \leq x + c$.

The corresponding indirect utility of Firm is

$$U(s) = \begin{cases} \mathbb{E}[\omega|s] & \text{if } \mathbb{E}[\omega|s] < x - c \\ x & \text{if } \mathbb{E}[\omega|s] \in [x - c, x + c] \\ \mathbb{E}[\omega|s] - c & \text{if } \mathbb{E}[\omega|s] > x + c. \end{cases}$$

A visual representation of Firm's indirect utility is in Figure 1b. Note that this utility depends only on posterior mean $\mathbb{E}[\omega|s]$, not on other aspects of posterior F_s. This implies that the analysis can focus on the distribution over posterior means instead of the distribution of posterior distributions. Without loss of generality, let s denote the mean of posterior distribution following realization s. Given prior F, there exists an experiment that produces a distribution G of posterior means if and only if F is a mean-preserving spread of G (Blackwell, 1953; Gentzkow and Kamenica, 2016; Kolotilin, 2018). I can thus rewrite Firm's optimal learning problem as

$$\max_{G \in \mathrm{MPC}(F)} \int U(s) dG(s), \tag{1}$$

where MPC(F) denotes the set of all mean-preserving contractions of F.

Below I restrict our attention to the case where Firm's status quo payoff is moderate $x \in (c, 1-c)$ and solve for Firm's optimal experiment. Cases where this condition fails to hold is trivial. For sufficiently small x, Regulator never goes to court and thus Firm perfectly learns ω and goes to court whenever ω is good enough. When both x and the prior mean of ω are sufficiently large, Firm conversely learns nothing and both players stay silent in equilibrium. A complete characterization of the equilibrium is in Appendix B.

3.2 Low-risk Case

I first consider Firm's ideal scenario. Intuitively, Firm would want to go to court when the underlying state is favorable and otherwise induce mutual silence.[4] To achieve this, Firm prefers to learn perfectly when the underlying state is good and not learn anything else—commission a study that reveals only whether ω is very high—as information generated about the bad states may encourage Regulator to go to court. Importantly, however, the absence of information would then imply that ω is *not* very high. For Firm to achieve her best outcome, she needs Regulator to still believe that ω is high *enough* even after knowing that ω is not very high in order to induce his silence; otherwise, Regulator updates his belief that ω is low and goes to court.

This can be achieved in the low-risk case where $\mathbb{E}[\omega \mid \omega < x + c] > x - c$. Here the prior belief of both players about ω is sufficiently high; even when ω is expected to not be very high ($\omega \leq x + c$), they believe that it is more likely to be moderate than high. Firm in this

[4]Note that the absolute first-best scenario for Firm would be to "trick" Regulator into believing that ω is low when in fact it is not, thereby persuading him to go to court. Firm essentially passes on the cost to Regulator while still revealing the underlying state. This, however, is never an optimal strategy for Firm within the framework of this model. For Firm to achieve this outcome, she would need to sufficiently pool the good states with the bad states to convince Regulator that ω is low. However, this is too costly for the bad states, as Regulator in equilibrium goes to court after the signal; Firm never prefers this from an ex-ante perspective. In Section 6.2, I discuss how this can be optimal for privately informed Firm.

case designs an experiment that perfectly separates all the good states ($\omega > x + c$) and pools the rest, which induces her ideal outcome. When the resulting signal from the experiment is a perfectly revealed $\omega > x + c$, Firm successfully goes to court in equilibrium and achieves a payoff higher than her initial status quo payoff. When players receive no information, both players know that $\omega \le x + c$, but the conditional posterior belief is still high enough that Regulator refrains from going to court and both players remain silent in equilibrium.

How do we verify that this is indeed the Firm-optimal experiment? Since ω is continuous, standard concavification method (Kamenica and Gentzkow, 2011) may have limitations when solving (1). I apply a novel technique recently developed by Dworczak and Martini (2019) to verify whether a candidate solution G is optimal. By this approach, I consider the smallest convex function ϕ that uniformly stays above U and identify the set of values such that $\phi(s) = U(s)$. Then, I see whether there is a distribution G such that $\text{supp}(G) \subseteq \{s : \phi(s) = U(s)\}$ and $\int \phi(s) dF(s) = \int \phi(s) dG(s)$. This allows for a simple graphical analysis to characterize Firm's optimal learning choice.

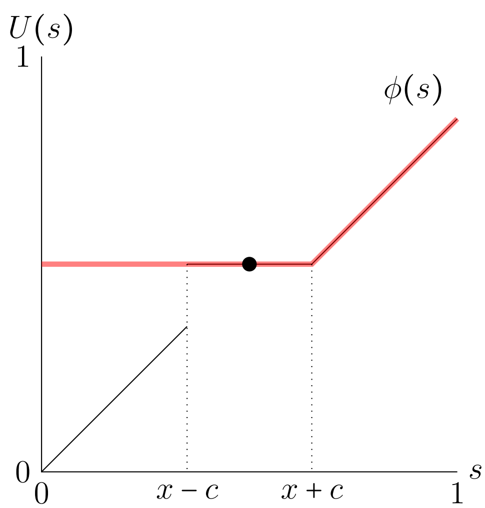

Figure 2: This figure plots Firm's indirect utility function $U(s)$ (black solid) and the smallest feasible convex function $\phi(s)$ (red translucent) above $U(s)$, given $c < x < 1 - c$. This can be used to apply the results by Dworczak and Martini (2019) for the low-risk case where $\mathbb{E}[\omega \mid \omega \le x + c] > x - c$.

Figure 2 illustrates Firm's indirect utility function $U(s)$ (black solid line) and the smallest feasible convex function $\phi(s)$ (red translucent line) above $U(s)$. For the low-risk case, the smallest convex function is a piece-wise linear function that bends at $x + c$. The ideal experiment described above is feasible and, therefore, optimal because its support $\{\mathbb{E}[\omega|\omega < x + c]\} \cup [x + c, 1]$ is a subset of [x-c,1]—the region on which $\phi = U$.[5]

Proposition 1 *If $\mathbb{E}[\omega \mid \omega \leq x + c] > x - c$, then it is optimal for Firm to fully learn $\omega > x + c$ and no other information.*

Proposition 1 summarizes the low-risk equilibrium. Firm here does not have to forego any loss of information while being able to perfectly deter Regulator from going to court and potentially revealing states that are detrimental to her.

3.3 Medium-risk Case

Now suppose that the previous case fails, and the product's risk is neither sufficiently low or sufficiently high; precisely, $\mathbb{E}[\omega|\omega \leq x + c] \leq x - c \leq \mathbb{E}[\omega]$. The product is expected to mandate only modest regulatory involvement, but the risk of court action is higher because conditional on the underlying state not being very high ($\omega \leq x+c$), it is likely to be extremely low. This is likely to be common in industries with polarized risk outcomes, where Firm's experiment may reveal that minimal regulation is sufficient due to significant societal benefits, or it might show the need for extremely stringent regulation because of the high risk of fatal accidents.

In such case, Firm-optimal still involves learning the best states, but now Firm is discouraged from learning all the good states under which it wants to go to court because with such design, Regulator goes to court whenever he learns that ω is not very high. To overcome this problem, Firm optimally pools some of the good states with the bad ones and only learns states larger than some cutoff $s^* > x + c$. This enables Firm to induce a conditional posterior

[5] In equilibrium, Firm goes to court when the underlying state is perfectly revealed via the experiment. In other words, we expect to observe Firm's action only when there is a "smoking gun." However, notice that there is a multiplicity of equilibria, where one of them involves only learning whether ω is below or above some threshold. Therefore, once we incorporate the costs of informative experiments where a more informative experiment is more costly, e.g., Matějka and McKay (2015) and Pomatto, Strack and Tamuz (2023), with sufficiently high cost we would expect Firm to commission such a study instead of the one provided in the main analysis.

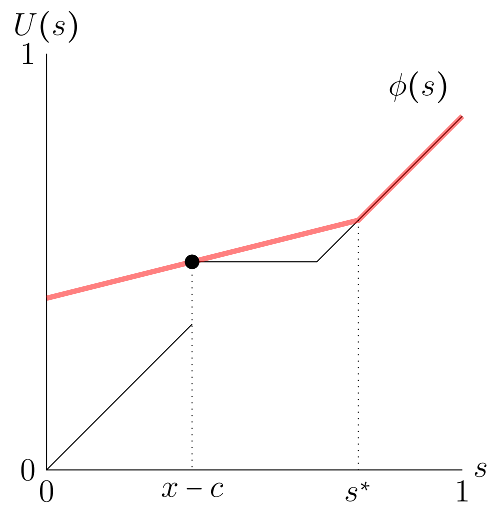

Figure 3: This figure represents the medium-risk case where $\mathbb{E}[\omega \mid \omega \leq x+c] \leq x-c \leq \mathbb{E}[\omega]$.

mean $\mathbb{E}[\omega|\omega \leq s^*]$ exactly at $x-c$, which is the minimum value that induces Regulator's silence. By only partially revealing high ω, Regulator believes that conditional on $\omega < s^*$—i.e., knowing that ω is not very high—there is still possibility that $\omega \in [x+c, s^*]$, and this shifts the posterior up to $x-c$ and deters him from going to court. This is illustrated in Figure 3. The smallest convex function is a piece-wise linear function that connects $(x-c)$ and (s^*, s^*). Cutoff s^* satisfies $\mathbb{E}[\omega|\omega \leq s^*] = x-c$. The optimal experiment G under moderate risk then has a support $\{\mathbb{E}[\omega|\omega \leq s^*]\} \cup (s^*, 1]$, where it coincides with prior distribution F for $\omega > s^*$, and the remaining probability mass is concentrated at $\omega = \mathbb{E}[\omega|\omega \leq s^*]$ as an atom.

Proposition 2 *If $\mathbb{E}[\omega|\omega \leq x+c] \leq x-c \leq \mathbb{E}[\omega]$, then it is optimal for Firm to fully learn $\omega > s^*$ and no other information, where s^* is the value such that $\mathbb{E}[\omega|\omega \leq s] = x-c$.*

Proposition 2 states the condition for the above equilibrium where Firm only partially learns the good states. Firm goes to court if $\omega > s^*$ and both players stay silent otherwise. Compared with Figure 2, $(x+c, s^*]$ represents the interval of information that the Firm has to "give up" in order to deter Regulator from going to court. Firm seeks to deter Regulator while minimizing the amount of information she forgoes by choosing s^* that satisfies $\mathbb{E}[\omega|\omega \leq s^*] = x-c$, which is the smallest value that convinces Regulator that going to court is not worth the risk.

3.4 High-risk Case

The last case involves both players believing that the product is likely to be harmful and require a high degree of regulatory control ($\mathbb{E}[\omega] < x - c$). A sufficiently small $\mathbb{E}[\omega]$ implies that the above strategy of pooling some good states with the bad states now cannot be optimal, as Regulator goes to court even if Firm gives up all the good states ($s^* = 1$). Now, in order to minimize Regulator's probability of going to court, Firm instead has to design an experiment that reveals whether the product is very *harmful*. By partially separating out the bad states, Firm maintains the conditional posterior of the less-bad states just high enough that deters Regulator from going to court.

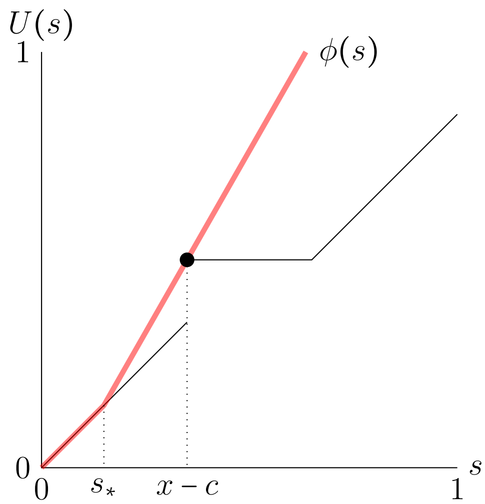

Figure 4: This figure represents the high-risk case where $\mathbb{E}[\omega] < x - c$.

This can be verified with the visual tool in Figure 4. If $\mathbb{E}[\omega] < x-c$, by Bayes plausibility, the smallest convex function must pass through point $(x-c, x)$, i.e., some states below $x-c$ must be pooled with those above. Then, $\phi(s)$ is a piece-wise linear function that connects (s_*, s_*) and $(x-c, x)$ where $E[\omega|\omega \geq s_*] = x - c$. The support of Firm's optimal choice of G is $\{\mathbb{E}[\omega|\omega \geq s_*]\} \cup [0, s_*)$, where it coincides with the prior distribution F for $\omega < s_*$, with the remaining probability mass concentrated at the point $\omega = \mathbb{E}[\omega|\omega \geq s_*]$ as an atom. This can be interpreted as a signal that fully reveals ω below s_* but provides no other information. Since $\mathbb{E}[\omega|\omega \geq s]$ increases in s, it follows that s_* that satisfies $E[\omega|\omega \geq s_*] = x - c$ is the smallest s that induces a posterior mean large enough to persuade Regulator. Firm's optimal choice then is to minimize this probability by revealing small enough ω below s_* and induce mutual silence otherwise.

Proposition 3 *If $\mathbb{E}[\omega] < x - c$, then it is optimal for Firm to fully learn $\omega < s_*$ and no other information, where s_* is the value such that $\mathbb{E}[\omega|\omega \geq s] = x - c$.*

Proposition 3 states the condition under which Firm fully learns small $\omega < s_*$ and lets Regulator go to court. Notably, if prior mean $\mathbb{E}[\omega]$ is sufficiently small, Firm *voluntarily* learns when the product is extremely harmful in order to successfully persuade Regulator not to go to court when the experiment returns a less-bad outcome. This suggests that empirically, we would expect to see firms disclosing damaging information, e.g., by commissioning a study or hiring an independent board of experts, specifically when the product faces a significant controversy.

4 Comparative Statics

In this section, I detail the implications of prior distribution F, Firm's status quo payoff x, and cost c on Firm's optimal choice of signal and her equilibrium utility.

Proposition 4 *Both s_* and s^* decrease as F increases in the sense of first-order stochastic dominance. Equivalently, the equilibrium informativeness of the experiment is non-monotonic with respect to her prior belief about the product.*

Proposition 4 establishes that the cutoffs used in the medium- and high-risk equilibrium both decrease as F increases in the sense of first-order stochastic dominance. Under medium-risk, Firm in equilibrium learns all states *above* s^* and no other information. A decrease in s^* then means that Firm's choice of experiment reveals more information in equilibrium. The idea is that as Firm's prior belief about ω increases, it takes a lower s^* to persuade Regulator not to go to court; Firm is thus able to learn more about the good states and exploit information while at the same time keeping Regulator silent. In the high-risk equilibrium, Firm uses cutoff s_* and learns states *below* it. A decrease in s_* implies that Firm learns less information. This is coming from a similar mechanism where it takes a lower s_* to induce

Regulator's silence, and so Firm can disclose less negative information about its product and still achieve deterrence. Note however that the prior belief about ω is still low that it is suboptimal for her to learn any good states. Firm therefore selects a strictly *less* informative experiment in this region.

Proposition 5 *Firm's equilibrium payoff is non-monotonic with respect to her status quo payoff x. In particular, Firm prefers lower x if $\mathbb{E}[\omega] < x - c$.*

Next, I examine Firm's utility with respect to two main parameters in the model, x and c. A notable finding from Proposition 5 is that Firm's equilibrium utility may decrease in her status quo payoff x. We observe this in the high-risk equilibrium, where the direct effect of an increase in x is outweighed by the indirect effect of an increase in Regulator's probability of going to court. Importantly, this implies that Firm in this region would prefer to *concede* some of her status quo payoff to Regulator. I discuss the implication of this result further in Section 6.1 by allowing players to bargain over their status quo division prior to playing the game and show that Pareto optimal divisions exist.

Similarly, the effect of cost c on Firm's utility is also not straightforward. Specifically, her utility decreases in c under low-risk and increases in c under high-risk; its effect under medium-risk is countervailing and will depend on the shape of prior F. Note that the non-monotonicity is mainly driven by the fact that c is a shared parameter for both players in the model. While an increase in c helps lower the probability of Regulator going to court, it simultaneously increases the cost of Firm challenging the status quo as well - the effect of c is therefore convoluted. In Appendix F I consider heterogeneous costs for players, and as expected, Firm's utility monotonically increases with the cost borne solely by Regulator.

5 Preference for Silence

So far, I have characterized information structures that maximize Firm's utility. Firm in equilibrium optimally manipulates the observability of the experiment in a way that induces Regulator's silence. By selectively revealing information about the product, the resulting posterior beliefs in medium- and high-risk equilibria are just high enough that dissuades

Regulator from going to court. In other words, Regulator in these cases remains silent in equilibrium when he could have received a higher payoff by going to court.

In this section, I further elaborate on this point and examine how bad Firm's control over information is for Regulator. If Regulator is potentially deterred from receiving a higher payoff with Firm's provision of information, would it be better for Regulator to shut down Firm's information generation in the first place? This would be equivalent to the case where no additional information is generated about the product, i.e., both players make their decision based on their prior belief about ω. While this is sometimes the case, I show that Regulator sometimes prefers to let Firm design the experiment. Notably, Regulator benefits from Firm's choice of information not because more information informs Regulator about when he can navigate the regulatory landscape in his favor; rather, he benefits because information leads to more *silence*.

Proposition 6 *Consider a low-risk case. Regulator prefers Firm's choice of information over no information if* $\mathbb{E}[\omega] > x + c$ *and* $\mathbb{E}[\omega|\omega < x + c] > x$.

Proposition 6 tells us that Regulator in the low-risk case has mixed preferences toward Firm's choice of experiment. In particular, he benefits from the information that Firm designs if $\mathbb{E}[\omega|\omega < x + c] > x$. First consider types $\omega \in [0, x)$. These types are always better off with no information; with no information, Regulator goes to court and successfully challenges the status quo, while with Firm's endogenous information Regulator is deterred from going to court and receives his status quo payoff which is always smaller than ω.

Types $\omega \in [x, x+c)$, however, are worse off with no information. To see why, note that we are considering a low-risk case; Firm thus goes to court in this region in the absence of any experiment, and Regulator's payoff is then always worse than his status quo payoff ($\omega < 1-x$). However, with Firm's endogenous choice of information, Firm only learns whether $\omega > x + c$ and stays silent otherwise. This means that Regulator gets to keep his status quo payoff $1 - x$, and he strictly prefers this over the equilibrium outcome given no information. Lastly, types $\omega \in [x + c, 1]$ are indifferent between the two options. From an ex-ante perspective, the costs of types $\omega \in [x, x + c)$ always outweigh the benefits of $\omega \in [0, x)$, and Regulator thus prefers to let Firm choose information in this region. Note that by definition Firm is always weakly better off choosing its own optimal experiment, and thus it is sometimes *mutually beneficial* to let Firm choose her preferred level of information.

This result in particular speaks to the existing literature that has emphasized how firms tend to control how and what information is revealed to regulatory agencies. This can easily produce information asymmetries that limit the regulator's ability to implement socially efficient regulatory policies (much like the seminal oversight points raised by Banks and Weingast (1992)). Furthermore, these asymmetries can generate equilibrium behaviors that have the appearance of agency capture, such as regulatory systematically favoring regulated firms (Pagliari, 2012; Shapiro, 2012; Carpenter and Moss, 2013; Beyers and Arras, 2020), or more specifically information capture, where the industry tries to manipulate the information on the basis of which the decision is made (Wagner, 2009; Agrell and Gautier, 2012).

This model finds that Regulator may optimally let this happen. He prefers Firm's choice of information despite the fact that it allows Firm to perfectly exploit information and go to court whenever she can successfully overturn the status quo, while at the same time deterring Regulator from doing so. As detailed above, this is largely driven by Regulator's incentive to avoid failed legal action and maintain the status quo.

6 Some Extensions

This section discusses three extensions to the model that incorporate additional aspects of the regulatory environment. The analysis so far has imposed no particular assumption on prior distribution F. For the extensions, I consider a simple example to illustrate the model in a more tractable setting. Specifically, I examine the case where the state of the world ω is a binary value of either 0 or 1, where $\omega = 1$ with probability $\pi \in (0,1)$. Bernoulli distribution is a limit case of f in which the distribution variance reaches the supremum of the set of feasible values of variance given the bounds of support and mean; random variable ω only takes extreme values 0 and 1 with atoms of size $1-\pi$ and π, respectively. Substantively, this may capture the binary nature of a formal action where only one player "wins" the fight.

The qualitative results of the model are similar with the baseline results in Section 3; the main difference is that Firm-optimal now involves partially learning either $\omega = 0$ or 1.[6] A complete analysis of the baseline model with this setup is in Appendix E.

6.1 Pre-game Settlement

First, I consider an extension of the model where Firm may offer to transfer resources to Regulator after observing cost c and status quo division $(x, 1-x)$ (initial stage) but prior to Firm's experimentation (signal stage). Specifically, Firm decides whether to keep the current status quo division $(x, 1-x)$ or propose a new division of dollar $(x', 1-x')$.[7] If Firm keeps the current division, players move on to play the benchmark game. If Firm proposes a new division, Regulator decides whether or not to accept the proposal. If Regulator accepts, players move on to play the game with the new division; if he rejects, they play the game with the initial division.

The settlement may be in the form of a side payment that shapes players' utilities from the status quo, or they could take a more legal form such as a certain solution (e.g., consent decrees involving additional warnings) proposed by Firm that Regulator may decide whether or not to agree on (McCarty, 2013). Note that I also allow Firm to demand compromise from Regulator ($x' > x$), which can be interpreted as Firm demanding a more lenient regulation from Regulator. I look for Firm's optimal offer $(x', 1-x')$ given status quo division $(x, 1-x)$, cost c, and probability π. As baked into the structure, if the new division is chosen in equilibrium, it will always Pareto dominate the status quo division. Otherwise, Firm will rather keep the current status quo division or Regulator will reject the offer.

Figure 5 visualizes the conditions under which a player is willing to compromise his or her status quo division prior to playing the game. More formally, each shaded region represents values of x and π where Pareto improving division $(x', 1-x')$ exists. Most notable in this

[6] Low-risk case does not exist in a binary setting. This is because if Firm perfectly learns all of its good state $\omega = 1$, in a binary setting it also perfectly reveals $\omega = 0$.

[7] Equivalently, Firm may offer Regulator payment $k \in [x-1, x]$, which would result in a new status quo division $(x-k, 1-x+k)$. When $k = 0$, Firm does not make an offer; when $k > 0$, she concedes some of her pie to Regulator; when $k < 0$, Regulator concedes some of his pie to Firm.

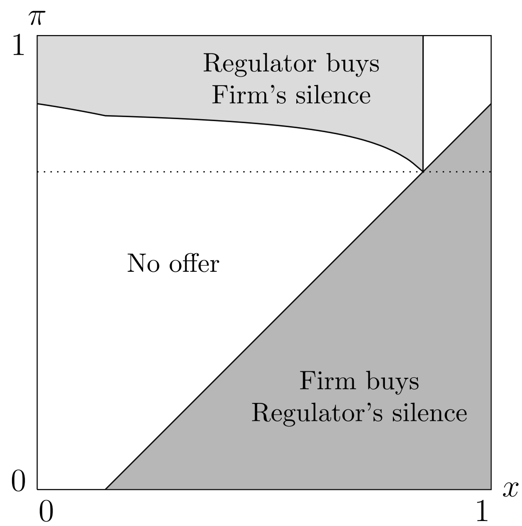

Figure 5: This figure illustrates the conditions under which a player is willing to compromise x to induce more silence from the other player. In the light-shaded region, Regulator buys Firm's silence; in the dark-shaded region, Firm buys Regulator's silence. I fix c at 0.15.

figure is that both Firm and Regulator sometimes prefer to compromise some of the initial share. Firm proposes to either take more ($x' > x$) or less ($x' < x$) share of the pie and Regulator accepts the offer. With a new division, both players become sufficiently satisfied that they are less incentivized to go to court, essentially "buying each other's silence." For Regulator, a settlement provides a reduced chance to end up in court and lose the case; for Firm, a settlement may help limit negative publicity, also avoiding the possibility of the penalty that would result if it loses (Krause, 2016).

The light-shaded region represents the case where Regulator buys Firm's silence. In equilibrium, Firm proposes to take more share of the pie ($x' > x$) and Regulator accepts the offer. Firm learns nothing afterwards, and both players simply decide to avoid pushing the issue to the forefront. By "offering" some portion of his initial share to Firm, Regulator gains an additional payoff by eliminating Firm's incentives to go to court. The dark-shaded region delineates the area where Firm conversely offers her share of the pie to Regulator. In the region above the dotted line in Figure 5, Firm's compromise leads to mutual silence where no legal action occurs in equilibrium; in other cases, Firm buys Regulator's silence by conceding x but now goes to court after signal 1 with the new division. Further note that when both players stay silent in equilibrium with the initial division, Firm never proposes a new division, i.e. there is no offer that Pareto dominates the original division. This means that bargaining over the status quo division strictly increases silence. The pre-game bargaining, in this sense, further allows players to keep the issue in the dark.

Proposition 7 *Suppose players can bargain over their initial status quo division.*

1. *Firm always prefers to give up her pie and buy Regulator's silence. The new division results in collusion if the prior probability is high ($\pi > 1 - 2c$) and Firm goes to court otherwise.*

2. *Regulator prefers to give up the pie in exchange for Firm's silence only with sufficiently high prior ($\pi > 1 - 2c$).*

Proposition 7 delineates conditions under which each player prefers to buy the other player's silence. Note that whenever Regulator goes to court with the initial division (dark-shaded region), Firm is willing to concede some of her status quo division to minimize this probability. Firm proposes a new division such that makes Regulator indifferent between accepting and rejecting the offer, and Regulator accepts in equilibrium.

With sufficiently high π, Regulator is willing to compromise his status quo division and offer a larger x' to discourage Firm from going to court (light-shaded region). Firm also prefers larger x in this range—she prefers to take a larger division and not go to court rather than to take the risk of a failed court action ($\omega = 0$) with a smaller division—and accepts the offer. Again, we observe that any successful bargain strictly results in increased silence, further emphasizing the point in Section 5 regarding players' preference for silence.

6.2 Privately Informed Firm

Consider an alternative setting where Firm chooses an experiment *after* observing ω. This is now an *informed information design* problem in the sense of Koessler and Skreta (2023). A central issue in this class of problems lies in examining the designer's interim incentives after learning additional information. My analysis in the benchmark model looked for Firm's ex-ante optimal signal. I check whether a privately informed Firm has the incentive to deviate from the proposed choice of signal and reveal her knowledge to Regulator. Firm's signal from the benchmark model is robust to private information in the low-risk equilibrium; in the medium- and high-risk equilibria, Firm's ex-ante optimal strategy unravels.

First, in the low-risk equilibrium, recall that Firm learns all the good states ($\omega > x + c$) and no other information. This means that Firm in equilibrium goes to court whenever it is ex-post optimal for her; privately observing ω thus does not change any of her strategy. Specifically, if Firm learns $\omega \geq x - c$, her payoff from the ex-ante optimal strategy is identical to the payoff she receives from disclosing her type. If $\omega < x - c$, Firm does better by not disclosing her type. Disclosing her type leads to Regulator taking action in equilibrium and leaves Firm with payoff ω, but Firm can do better by following the ex-ante optimal strategy and receiving x. Thus, Firm can never do better by deviating and disclosing her type.

Next, consider the medium-risk equilibrium where Firm only partially learns the good states ($\omega > s^*$) and pools the rest and suppose that Firm privately learns state ω. If $\omega < x - c$, Firm again does better with the ex-ante optimal strategy than disclosing her type, as she can induce Regulator's silence despite the product being socially harmful. If $\omega \in [x - c, x + c]$ or $\omega > s^*$, her payoff is identical with and without disclosure. However, if Firm learns that $\omega \in (x + c, s^*]$, she can receive payoff $\omega - c$ (that is larger than her payoff from the ex-ante optimal strategy, x) by disclosing information. It follows that these types will separate themselves, resulting in unraveling over the interval. Notice that this is the exact interval that Firm was willing to "give up" ex-ante in order to deter Regulator from going to court.

With high risk, the ex-ante optimal signal for Firm is to learn $\omega < s_*$ and let Regulator go to court for these cases; she induces mutual silence otherwise and yields a payoff of x. By pooling states $\omega > s_*$, Firm is essentially giving up her own opportunity to go to court ($\omega > x + c$) so that she can deter Regulator from going to court given $\omega \in [s_*, x - c)$. Similar to the logic above, if Firm learns that $\omega \in (x + c, 1]$, she can receive payoff $\omega - c$ larger than x by disclosing that information. Firm cannot credibly commit to making this sacrifice if she learns that her type is $\omega > x + c$, and so the unraveling logic also applies here.

This analysis clearly demonstrates Firm's trade-off when deterring Regulator from going to court. Whenever Firm ex-ante foregoes some opportunity to induce Regulator's silence, if she realizes that the state is favorable enough and that she can gain more by going to court, Firm's strategy in the benchmark model unravels. On the other hand, Firm's interim incentive constraints are satisfied when Firm is ex-ante already enjoying informational advantage at its full potential.

Interim optimal design. Then, what is the alternative strategy that is interim optimal for Firm when the ex-ante optimal strategy unravels with Firm's private information? In Appendix E.1, I fully characterize the interim optimal design for binary state ω. With private information, we now see an incentive of Firm to *persuade* Regulator to go to court. Suppose $\omega = 1$ and Firm learns this privately. Firm wants to go to court, but in an ideal scenario, Regulator would be the one who goes to court because Firm can then overturn the status quo without incurring the cost of action c. Accordingly, when the prior belief about ω is sufficiently low—and Regulator goes to court without additional information—it is interim optimal for Firm to provide no information.

6.3 "Getting It Right"

So far I have focused on a particular regulatory context where Firm and Regulator have misaligned preferences regarding the right level of regulation for the product. Realistically, however, there can be intervals where Firm and Regulator are aligned in their preferences, also preferring a more lenient regulation when he believes that the benefit of the product clearly outweighs the potential risks. Here I consider a different payoff matrix that reflects Regulator's preferences to "get it right" and show how the results of the benchmark model effectively captures all the key qualitative outcome, even without the added complexity of the extended version. The new payoff matrix is as follows:

Firm \ Regulator	N	C
N	$(x, -\lvert x-\omega\rvert)$	$(\omega, -c)$
C	$(\omega - c, 0)$	$(\omega - c, -c)$

Table 2: New Payoff Structure

Note that Regulator is now much more likely to go to court than in the benchmark model, as he simply benefits from *revealing* the underlying state and matching the regulatory standard to that state (unless the cost of doing so is too large).

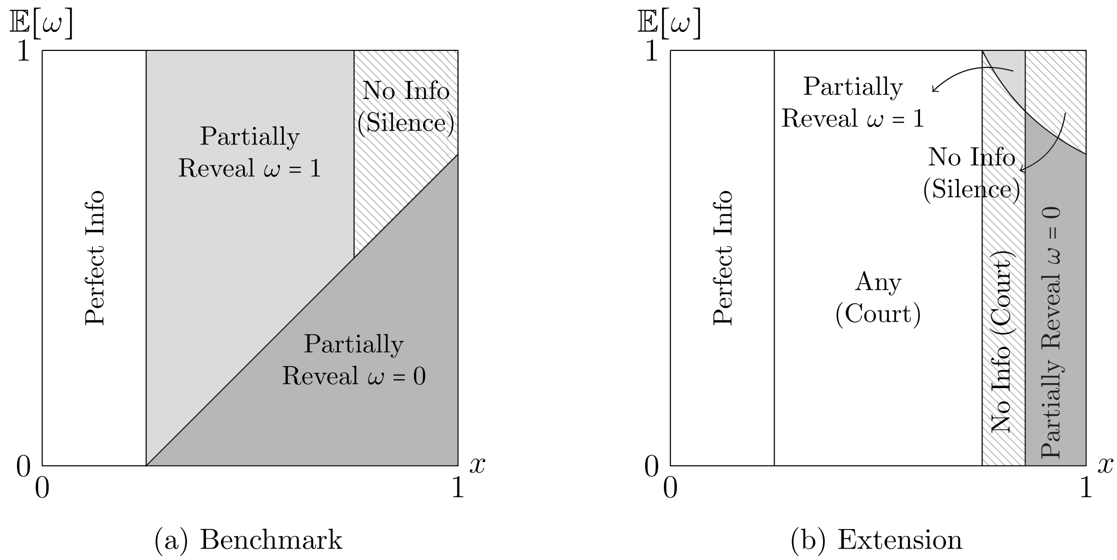

Figure 6: These figures plot how Firm's optimal choice of signal depends on status quo x and prior mean $\mathbb{E}[\omega]$ given $c = 0.25$. The left panel represents the benchmark model, and the right panel represents the extended model.

Figure 6 represents a visual comparison of the results from the benchmark and the extended model. Note that Figure 6a represents the results of the benchmark model with a binary setup. In Figure 6b, the same set of strategies as the benchmark model is optimal for Firm, but notice that any and no information is more frequently an optimal signal strategy for Firm. When x is moderate, Firm is indifferent to any choice of experiment because the underlying state $\omega \in \{0, 1\}$ is sufficiently distant from the status quo that Regulator is willing to pay the cost and reveal ω regardless of the signal he received via Firm's experiment; thus, persuasion becomes irrelevant. With higher x and relatively lower $\mathbb{E}[\omega]$, Regulator goes to court without any information, and interestingly, Firm intentionally encourages Regulator do so by providing no information - Firm takes some risk of ω turning out to be 0 but in return is able to reveal $\omega = 1$ without having to pay cost c.

7 Conclusion

To what degree would a firm seek to acquire information about an outcome when it intends to leverage that information for future action but what it learns is also available to a regulator with misaligned interest? I develop a model that examines this informational problem in the context of regulatory uncertainty and show that the firm-optimal can be achieved by designing a binary experiment. The experiment is informative only about whether the firm or the regulator has an incentive to invest in challenging the current level of regulation. Specifically, the firm may choose two very different forms of experiment: when it believes that the product requires only minimal regulation, the firm optimally commissions a study that only reveals whether the product is sufficiently beneficial. The firm in equilibrium goes to court and challenges the status quo whenever the experiment reveals good information. Even when the experiment informs players that the product is not beneficial, the resulting belief is still high enough that the regulator believes going to court is not worth the risk and stays silent. Conversely, when it believes the product demands extensive regulatory scrutiny, the firm designs an experiment that only reveals whether the product is sufficiently harmful. The firm intentionally commits to a design of experiment that reveals whenever a product necessitates extensive oversight. By doing so, the firm persuades the regulator to stay silent when the experiment informs players otherwise.

One avenue for future research is to understand how this dynamic applies to different institutional contexts that may potentially limit the firm's ability to persuade. While the regulator may be dependent on the firm for information, he (or a policymaker) can choose whether to grant discretion to an agent (Montagnes and Wolton, 2017) or institute requirements for information provision that may constrain the firm's ability to manipulate information (Doval and Skreta, 2024). Examining the firm's strategic incentives in these contexts could be useful.

References

Agrell, Per J and Axel Gautier. 2012. Rethinking regulatory capture. In *Recent advances in the analysis of competition policy and regulation*. Edward Elgar Publishing.

Austen-Smith, David. 1993. "Information and influence: Lobbying for agendas and votes." *American Journal of Political Science* 37(3):799–833.

Awad, Emiel. 2020. "Persuasive lobbying with allied legislators." *American Journal of Political Science* 64(4):938–951.

Banks, Jeffrey S and Barry R Weingast. 1992. "The political control of bureaucracies under asymmetric information." *American Journal of Political Science* pp. 509–524.

Beyers, Jan and Sarah Arras. 2020. "Who feeds information to regulators? Stakeholder diversity in European Union regulatory agency consultations." *Journal of Public Policy* 40(4):573–598.

Bils, Peter, Robert J Carroll and Lawrence S Rothenberg. 2020. Strategic Avoidance and Rulemaking Procedures. Technical report Working Paper. URL: https://www. wallis. rochester. edu/assets/pdf . . . .

Blackwell, David. 1953. "Equivalent comparisons of experiments." *The annals of mathematical statistics* pp. 265–272.

Boehmke, Frederick J, Sean Gailmard, John Wiggs Patty et al. 2006. "Whose ear to bend? Information sources and venue choice in policy-making." *Quarterly Journal of Political Science* 1(2):139–169.

Callander, Steven. 2011*a*. "Searching and learning by trial and error." *American Economic Review* 101(6):2277–2308.

Callander, Steven. 2011*b*. "Searching for good policies." *American Political Science Review* 105(4):643–662.

Callander, Steven and Bård Harstad. 2015. "Experimentation in federal systems." *The Quarterly Journal of Economics* 130(2):951–1002.

Callander, Steven, Dana Foarta and Takuo Sugaya. 2022*a*. "The dynamics of a policy outcome: Market response and bureaucratic enforcement of a policy change." *American Journal of Political Science* .

Callander, Steven, Dana Foarta and Takuo Sugaya. 2022*b*. "Market competition and political influence: An integrated approach." *Econometrica* 90(6):2723–2753.

Callander, Steven and Patrick Hummel. 2014. "Preemptive policy experimentation." *Econometrica* 82(4):1509–1528.

Carpenter, Daniel and David A Moss. 2013. *Preventing regulatory capture: Special interest influence and how to limit it*. Cambridge University Press.

Carpenter, Daniel, Justin Grimmer and Eric Lomazoff. 2010. "Approval regulation and endogenous consumer confidence: Theory and analogies to licensing, safety, and financial regulation." *Regulation & Governance* 4(4):383–407.

Carpenter, Daniel and Michael M Ting. 2007. "Regulatory errors with endogenous agendas." *American Journal of Political Science* 51(4):835–852.

Carpenter, Daniel P. 2002. "Groups, the media, agency waiting costs, and FDA drug approval." *American Journal of Political Science* pp. 490–505.

Carpenter, Daniel P. 2004. "Protection without capture: Product approval by a politically responsive, learning regulator." *American Political Science Review* 98(4):613–631.

Dellis, Arnaud. 2023. "Legislative informational lobbying." *Journal of Economic Theory* 208:105595.

Doval, Laura and Vasiliki Skreta. 2024. "Constrained information design." *Mathematics of Operations Research* 49(1):78–106.

Dworczak, Piotr and Giorgio Martini. 2019. "The simple economics of optimal persuasion." *Journal of Political Economy* 127(5):1993–2048.

Ellis, Christopher J and Thomas Groll. 2020. "Strategic legislative subsidies: Informational lobbying and the cost of policy." *American Political Science Review* 114(1):179–205.

Gailmard, Sean and John W Patty. 2012. *Learning while governing: Expertise and accountability in the executive branch.* University of Chicago Press.

Gailmard, Sean and John W Patty. 2017. "Participation, process and policy: the informational value of politicised judicial review." *Journal of Public Policy* 37(3):233–260.

Gentzkow, Matthew and Emir Kamenica. 2016. "A Rothschild-Stiglitz approach to Bayesian persuasion." *American Economic Review* 106(5):597–601.

Henry, Emeric, Marco Loseto and Marco Ottaviani. 2022. "Regulation with experimentation: Ex ante approval, ex post withdrawal, and liability." *Management Science* 68(7):5330–5347.

Henry, Emeric and Marco Ottaviani. 2019. "Research and the approval process: The organization of persuasion." *American Economic Review* 109(3):911–955.

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

Koessler, Frédéric and Vasiliki Skreta. 2023. "Informed information design." *Journal of Political Economy* 131(11):3186–3232.

Kolotilin, Anton. 2018. "Optimal information disclosure: A linear programming approach." *Theoretical Economics* 13(2):607–635.

Krause, Joan H. 2016. "Truth, Falsity, and Fraud: Off-Label Drug Settlements and the Future of the Civil False Claims Act." *Food & Drug LJ* 71:401.

Libgober, Brian D. 2020. "Strategic proposals, endogenous comments, and bias in rulemaking." *The Journal of Politics* 82(2):642–656.

Matějka, Filip and Alisdair McKay. 2015. "Rational inattention to discrete choices: A new foundation for the multinomial logit model." *American Economic Review* 105(1):272–298.

McCarty, Nolan. 2013. "Complexity, capacity, and capture." *Preventing regulatory capture: Special interest influence and how to limit it* pp. 99–123.

McCarty, Nolan. 2017. "The regulation and self-regulation of a complex industry." *The Journal of Politics* 79(4):1220–1236.

McClellan, Andrew. 2022. "Experimentation and approval mechanisms." *Econometrica* 90(5):2215–2247.

Montagnes, B Pablo and Stephane Wolton. 2017. "Rule versus discretion: regulatory uncertainty, firm investment, and bureaucratic organization." *The Journal of Politics* 79(2):457–472.

NTSB. 2013. Aircraft Incident Report: Auxiliary Power Unit Battery Fire, Japan Airlines Boeing 787-8, JA829J, Boston, Massachusetts, 7 January 2013. Technical report.

Pagliari, Stefano. 2012. "How can we mitigate capture in financial regulation?".

Patty, John W and Elizabeth Maggie Penn. 2022. "Ban The Box? Information, Incentives, and Statistical Discrimination." *arXiv preprint arXiv:2208.08348* .

Patty, John W and Elizabeth Maggie Penn. 2023. "Algorithmic fairness and statistical discrimination." *Philosophy Compass* 18(1):e12891.

Pomatto, Luciano, Philipp Strack and Omer Tamuz. 2023. "The cost of information: The case of constant marginal costs." *American Economic Review* 113(5):1360–1393.

Potters, Jan and Frans Van Winden. 1992. "Lobbying and asymmetric information." *Public Choice* 74:269–292.

Prato, Carlo and Ian R Turner. 2022. "Institutional foundations of the power to persuade." *Center for Open Science SocArXiv* (4w9af).

Sabel, Charles, Gary Herrigel and Peer Hull Kristensen. 2018. "Regulation under uncertainty: The coevolution of industry and regulation." *Regulation & Governance* 12(3):371–394.

Schnakenberg, Keith. 2017. "Informational lobbying and legislative voting." *American Journal of Political Science* 61(1):129–145.

Schnakenberg, Keith and Ian Turner. 2019. "Signaling with reform: How the threat of corruption prevents informed policy-making." *American Political Science Review* 113(3):762–777.

Schnakenberg, Keith and Ian Turner. 2021. "Helping friends or influencing foes: Electoral and policy effects of campaign finance contributions." *American Journal of Political Science* 65(1):88–100.

Schnakenberg, Keith and Ian Turner. 2023. "Formal Theories of Special Interest Influence." *Working Paper* .

Shapiro, Sidney A. 2012. "The complexity of regulatory capture: Diagnosis, causality, and remediation." *Roger Williams UL Rev.* 17:221.

Wagner, Wendy E. 2009. "Administrative law, filter failure, and information capture." *Duke LJ* 59:1321.

Appendix for

Strategic Experiments Under Regulatory Uncertainty

Contents

A Proof

Proposition 1 *If* $\mathbb{E}[\omega \mid \omega \leq x + c] > x - c$*, then it is optimal for Firm to fully learn* $\omega > x + c$ *and no other information.*

Proof: If the experiment reveals that $\omega > x + c$, the posterior mean is $s = \omega$. Since $s > x + c$, Firm chooses to go to court, resulting in a payoff of $s - c$. Conversely, if $\omega \leq x + c$, the posterior mean is pooled, yielding $s = \mathbb{E}[\omega \mid \omega \leq x + c]$. Since $s \in (x - c, x + c]$, neither goes to court, and Firm receives a payoff of x. Graphically, we can verify that the induced posterior means and the corresponding payoffs lie at the intersection of the indirect utility function $U(s)$ and the smallest feasible convex function $\phi(s)$:

$$\phi(s) = \begin{cases} x & \text{if } s \leq x + c \\ s - c & \text{if } s > x + c \end{cases}$$

which is a piecewise linear function that passes through the points $(0, x)$, $(x + c, x)$, and $(1, 1 - c)$. ∎

Proposition 2 *If* $\mathbb{E}[\omega|\omega \leq x + c] \leq x - c \leq \mathbb{E}[\omega]$*, then it is optimal for Firm to fully learn* $\omega > s^*$ *and no other information, where* s^* *is the value such that* $\mathbb{E}[\omega|\omega \leq s] = x - c$*.*

Proof: If the experiment reveals that $\omega > s^*$, the posterior mean is $s = \omega$. Since $s > x+c$, Firm chooses to go to court, resulting in a payoff of $s-c$. Conversely, if $\omega \leq s^*$, the posterior mean is pooled, yielding $s = x-c$. Since $s = x-c$, neither goes to court, and Firm receives a payoff of x. Graphically, we can verify that the induced posterior means and the corresponding payoffs lie at the intersection of the indirect utility function $U(s)$ and the smallest feasible convex function $\phi(s)$:

$$\phi(s) = \begin{cases} \frac{(s^* - x - c)s + (s^* + x - c)c}{s^* - x + c} & \text{if } s \leq s^* \\ s - c & \text{if } s > s^* \end{cases}$$

which is a piecewise linear function that passes through the points $(x-c, x)$, $(s^*, s^* - c)$, and $(1, 1 - c)$. ∎

Proposition 3 *If $\mathbb{E}[\omega] < x - c$, then it is optimal for Firm to fully learn $\omega < s_*$ and no other information, where s_* is the value such that $\mathbb{E}[\omega | \omega \geq s] = x - c$.*

Proof: If the experiment reveals that $\omega < s^*$, the posterior mean is $s = \omega$. Since $s < x - c$, Regulator chooses to go to court, and Firm receives a payoff of s. Conversely, if $\omega \geq s^*$, the posterior mean is pooled, yielding $s = x - c$. Since $s = x - c$, neither goes to court, and Firm receives a payoff of x. Graphically, we can verify that the induced posterior means and the corresponding payoffs lie at the intersection of the indirect utility function $U(s)$ and the smallest feasible convex function $\phi(s)$:

$$\phi(s) = \begin{cases} s & \text{if } s < s^* \\ \frac{(x - s^*)s - s^* c}{x - c - s^*} & \text{if } s \geq s^* \end{cases}$$

which is a piecewise linear function that passes through the points $(0, 0)$, (s^*, s^*), and $(x - c, x)$. ∎

Proposition 4 *Both s_* and s^* decrease as F increases in the sense of first-order stochastic dominance. Equivalently, the equilibrium informativeness of the experiment is non-monotonic with respect to her prior belief about the product.*

Proof: I examine the effect of F on equilibrium cutoffs. First, consider cutoff s_*, which satisfies

$$\mathbb{E}[\omega \mid \omega \geq s_*] = x - c.$$

All else equal (including s_*), the right-hand side of the equation increases as F increases in the sense of first-order stochastic dominance. Thus, s_* must decrease for the equation to hold. In the high-risk equilibrium, information is revealed with probability $F(s_*)$. Consider F' that first-order stochastically dominates F. Let s_{**} denote the cutoff in equilibrium given F'. Comparing $F(s_*)$ and $F'(s_{**})$, it follows that $F(s_*) > F'(s_*)$ by the definition of first-order stochastic dominance, and $F'(s_*) > F'(s_{**})$ because the cutoff decreases in F. Therefore, the probability that some information is revealed decreases in F for the high-risk case. At the "upper bound" where $\mathbb{E}[\omega] \to x - c$, the probability reaches zero because $s_* \to 0$.

Second, consider cutoff s^*, which satisfies

$$\mathbb{E}[\omega \mid \omega \leq s^*] = x - c.$$

Similarly, all else equal, including s^*, the right-hand side of the equation increases as F increases in the sense of first-order stochastic dominance. Thus, s^* must decrease for the equation to hold. In the medium-risk equilibrium, information is revealed with probability $1 - F(s^*)$. Similarly, consider first-order stochastically dominant distribution F' and corresponding cutoff s^{**}. Comparing $1-F(s^*)$ and $1-F'(s^{**})$, it follows that $1-F(s^*) < 1-F'(s^*)$ by the definition of first-order stochastic dominance, and $1-F'(s^*) < 1-F'(s^{**})$ because the cutoff decreases in F. Therefore, the probability that some information is revealed increases in F in the medium-risk case. A similar logic applies to the low-risk case, where the cutoff is constant at $x + c$.

To summarize, as F increases in the sense of first-order stochastic dominance, the equilibrium informativeness of the experiment initially decreases, eventually reaching zero in the high-risk case. Subsequently, the informativeness increases in the medium- and low-risk cases. ∎

Proposition 5 *Firm's equilibrium payoff is non-monotonic with respect to her status quo payoff x. In particular, Firm prefers lower x if $\mathbb{E}[\omega] < x - c$.*

Proof: In the low-risk equilibrium, Firm receives

$$\int_0^{x+c} (x) dF(\omega) + \int_{x+c}^{1} (\omega - c) dF(\omega).$$

Consider a larger $x' > x$ such that results in low-risk case equilibrium. Firm receives

$$\begin{aligned}
&\int_0^{x'+c}(x')dF(\omega) + \int_{x'+c}^1(\omega - c)dF(\omega) \\
=&\int_0^{x+c}(x')dF(\omega) + \int_{x+c}^{x'+c}(x')dF(\omega) + \int_{x'+c}^1(\omega - c)dF(\omega) \\
>&\int_0^{x+c}(x)dF(\omega) + \int_{x+c}^{x'+c}(\omega - c)dF(\omega) + \int_{x'+c}^1(\omega - c)dF(\omega) \\
=&\int_0^{x+c}(x)dF(\omega) + \int_{x+c}^1(\omega - c)dF(\omega),
\end{aligned}$$

which is larger than the payoff induced by the original x. In other words, conditional on the low-risk equilibrium, Firm's utility increases in x.

In the medium-risk equilibrium, Firm receives

$$\begin{aligned}
&\int_0^{s^*}(x)dF(\omega) + \int_{s^*}^1(\omega - c)dF(\omega) \\
=&\int_0^{s^*}(\omega + c)dF(\omega) + \int_{s^*}^1(\omega - c)dF(\omega) \\
=&\int_0^1(\omega)dF(\omega) + 2\int_0^{s^*}(c)dF(\omega) - c.
\end{aligned}$$

Note that s^* increases in x because the right-hand side of the condition $\mathbb{E}[\omega|\omega \leq s^*] = x - c$ increases in x. Then, conditional on the medium-risk equilibrium, Firm's utility increases in x.

In the high-risk equilibrium, Firm receives

$$\begin{aligned}
&\int_0^{s_*}(\omega)dF(\omega) + \int_{s_*}^1(x)dF(\omega) \\
=&\int_0^{s_*}(\omega)dF(\omega) + \int_{s_*}^1(\omega + c)dF(\omega) \\
=&\int_0^1(\omega)dF(\omega) + \int_{s_*}^1(c)dF(\omega).
\end{aligned}$$

Note that s_* increases in x following a similar logic. Conditional on the high-risk equilibrium, Firm's utility decreases in x.

To summarize, Firm's utility is non-monotonic in x: it increases in x in the low- and medium-risk equilibria but decreases in x in high-risk equilibrium. Consequently, under the high-risk equilibrium, Firm prefers to begin the game with a lower x, if feasible. ∎

Proposition 6 *Consider a low-risk case. Regulator prefers Firm's choice of information over no information if* $\mathbb{E}[\omega] > x + c$ *and* $\mathbb{E}[\omega | \omega < x + c] > x$.

Proof: In the low-risk equilibrium, Regulator receives

$$\int_0^{x+c} (1 - x) dF(\omega) + \int_{x+c}^1 (1 - \omega) dF(\omega).$$

Without information, Firm goes to court because $\mathbb{E}[\omega] > x + c$. Then, Regulator receives

$$\begin{aligned} & \int_0^1 (1 - \omega) dF(\omega) \\ = & \int_0^{x+c} (1 - \omega) dF(\omega) + \int_{x+c}^1 (1 - \omega) dF(\omega) \\ < & \int_0^{x+c} (1 - x) dF(\omega) + \int_{x+c}^1 (1 - \omega) dF(\omega) \\ \Leftrightarrow & \int_0^{x+c} (1 - \omega) dF(\omega) < \int_0^{x+c} (1 - x) dF(\omega) \\ \Leftrightarrow & \int_0^{x+c} (x) dF(\omega) < \int_0^{x+c} (\omega) dF(\omega), \end{aligned}$$

which is equivalent to the assumption $\mathbb{E}[\omega | \omega < x + c] > x$. In other words, for the low-risk case, Regulator prefers Firm's choice of information over no information if (1) Firm goes to court without information and (2) the prior distribution conditional on $\omega < x + c$ is skewed to the left of x. ∎

Proposition 7 *Suppose players can bargain over their initial status quo division.*

1. *Firm always prefers to give up her pie and buy Regulator's silence. The new division results in collusion if the prior probability is high* ($\pi > 1 - 2c$) *and Firm goes to court otherwise.*

2. *Regulator prefers to give up the pie in exchange for Firm's silence only with sufficiently high prior* ($\pi > 1 - 2c$).

Proof: In this proposition, I only consider an example of binary ω. The complete analysis of this version of the model is in Appendix E. For the pre-game settlement to be feasible, there must exist Pareto improving x'. I first identify the set of such x'. Then, I look for Firm's optimal choice of x' if it exists.

1. If Firm goes to court in equilibrium given the initial x, there does not exist Pareto improving x'.

2. If Firm fully reveals ω in equilibrium given the initial x, there may exist Pareto improving x' such that induces mutual silence. This is because Regulator's payoff discontinuously jumps when the equilibrium transitions from Firm going to court to mutual silence. Firm's optimal choice of new division is

$$x' = \pi + (1 - \pi)x > 1 - c.$$

3. If Firm partially reveals $\omega = 1$ in equilibrium, there may exist Pareto improving x' such that induces mutual silence. Similarly, this is because Regulator's payoff discontinuously jumps when the equilibrium outcome transitions from Firm going to court to mutual silence. Firm's optimal choice of new division is

$$x' = \frac{c + \pi(1 - x)}{1 - x + c} \in (1 - c, c + \pi).$$

4. If Firm partially reveals $\omega = 0$ in equilibrium, there always exists Pareto improving x' such that induces mutual silence if $c > (1 - \pi)/2$ and Firm going to court otherwise. This is because Regulator's payoff is constant in x while Firm's payoff is decreasing in x. Firm's optimal choice of new division is

$$x' = c + \pi.$$

To summarize, if Regulator goes to court in equilibrium (after Firm partially reveals $\omega = 0$), then Firm always has the incentive to offer $x' = c + \pi$ which is smaller than the initial division $x > c + \pi$. If Firm goes to court in equilibrium (after Firm fully or partially reveals $\omega = 1$), then Regulator may have the incentive to accept a new offer $x' > x$ such that induces mutual silence. The necessary condition for such an offer to exist is $\pi > 1 - 2c$, as implied by the conditions for x' above. ∎

B Complete Characterization of Equilibrium

In this section, I provide a complete equilibrium analysis of the benchmark model. As shown in Figure B.1, $U(s)$ can take largely four different shapes depending on x and c. First, Figure B.1a shows that $U(s)$ is constant at x when $x \in (1-c, c)$. Firm's choice of G does not matter here; any s will lead to a subgame equilibrium where both remain silent. Next, Figure B.1b shows that $U(s)$ is already continuous and convex when $x \le \min\{c, 1-c\}$. This implies that Firm's optimal choice here is to generate perfect information.

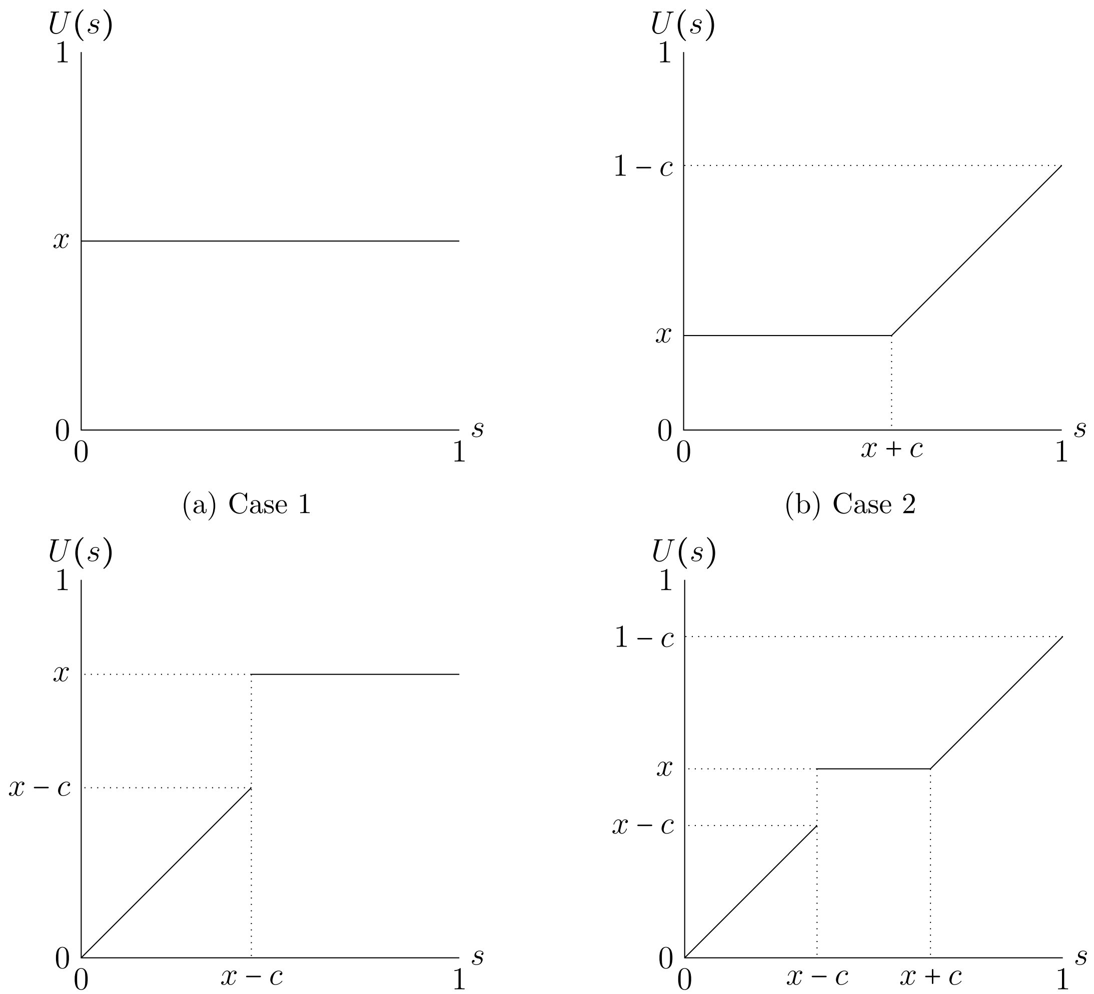

(d) Case 4

Figure B.1: This figure plots Firm's indirect utility function $U(s)$ of which shape depends on x, c. The first case is when $x \in (1-c, c)$. The second case is when $x \le \min\{c, 1-c\}$. The third case is when $x \ge \max\{c, 1-c\}$. The fourth case is when $x \in (c, 1-c)$.

The remaining two cases require further analysis because the indirect utility is not continuous. Figure B.2 divides Figure B.1c into two subcases. Figure B.2a is when the risk is low for Firm, i.e., $\mathbb{E}[\omega] \geq x - c$. Without information, both players remain silent and this is the best outcome for Firm. Therefore, Firm's optimal choice is to generate no information. Figure B.2b is when the risk is high for Firm, i.e., $\mathbb{E}[\omega] < x - c$. Since Regulator goes to court without information, Firm pools ω above a cutoff to deter going to court. It follows that the optimal cutoff is s_* such that satisfies $\mathbb{E}[\omega|\omega \geq s_*] = x - c$, which is the minimal cutoff that deters Regulator going to court.

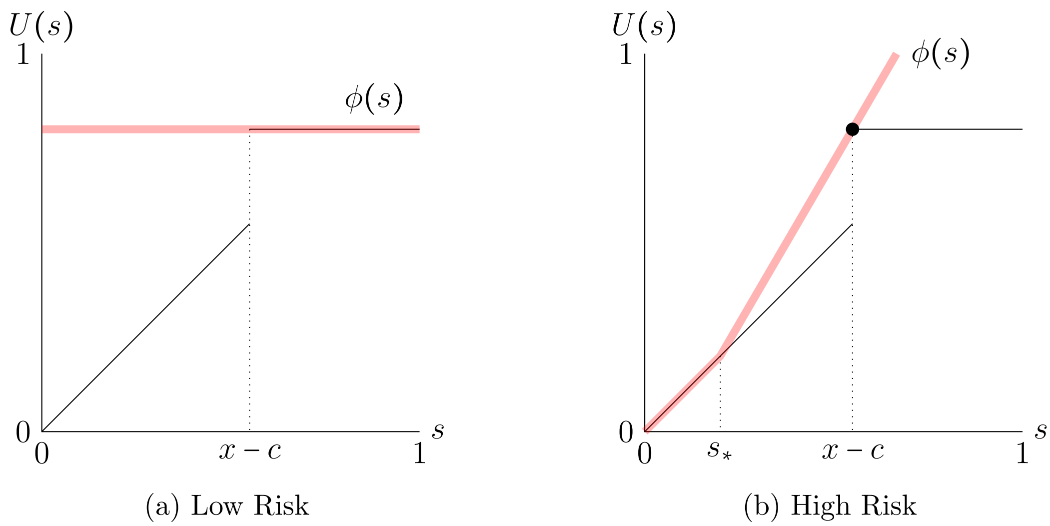

Figure B.2: This figure plots Firm's indirect utility function $U(s)$ (black solid) and the smallest feasible convex function $\phi(s)$ (red translucent) above $U(s)$, given $x \geq \max\{c, 1 - c\}$. The left panel represents the case where $\mathbb{E}[\omega] < x - c$. The right panel represents the case where $\mathbb{E}[\omega] \geq x - c$.

Figure 2-4 in the main section are three subcases of Figure B.1d. Figure 2 is when the risk is low for Firm, i.e., $\mathbb{E}[\omega|\omega \leq x + c] > x - c$. As shown by $\phi(s)$, Firm reveals ω above $x + c$ and goes to court; both players remain silent otherwise. Figure 3 is when the risk is medium for Firm, i.e., $\mathbb{E}[\omega|\omega \leq x + c] \leq x - c \leq \mathbb{E}[\omega]$. Here, Firm cannot induce mutual silence by pooling ω below $x + c$. It follows that Firm must employ a larger cutoff to do so. The optimal cutoff is then s^* such that satisfies $\mathbb{E}[\omega|\omega \leq s^*] = x - c$, which is the minimal cutoff that can deter going to court. Figure 4 is when the risk is high for Firm, i.e., $\mathbb{E}[\omega] < x - c$. This is equivalent to that from Figure B.2b.

Figure B.3 illustrates how the equilibrium outcome depends on x and $\mathbb{E}[\omega]$. For suffi-

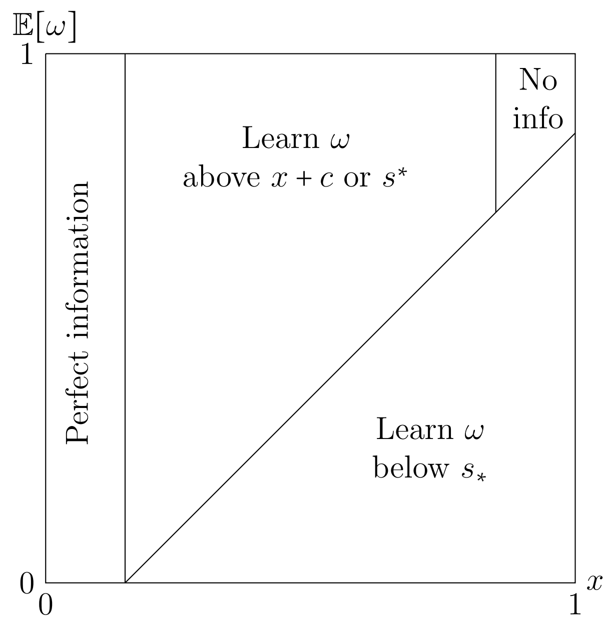

Figure B.3: This figure plots how Firm's optimal choice of the signal depends on status quo x and prior mean $\mathbb{E}[\omega]$ given $c = 0.15$.

ciently small x, Firm fully reveals ω because Regulator does not have the incentive to go to court for all beliefs. Then, Firm can perfectly learn ω without any consequence because information cannot be used against her. For sufficiently large x and $\mathbb{E}[\omega]$, Firm learns nothing. This is because Firm is already satisfied with the status quo, and Regulator has no incentive to go to court without information. For x relatively larger than $\mathbb{E}[\omega]$, Firm reveals ω below s_*; this is when Regulator is confident enough to go to court without information. Therefore, Firm gives up the worst states in order to induce mutual silence otherwise. For moderate x and sufficiently large $\mathbb{E}[\omega]$, Firm reveals ω above $x + c$ or s^*; the exact cutoff is determined by the condition $\mathbb{E}[\omega|\omega \leq x + c] \leq x - c$. To summarize,

1. Firm's choice of signal does not matter if $x \in (1 - c, c)$ because neither goes to court for all beliefs.

2. It is optimal for Firm to fully learn ω if $x \leq \min\{c, 1 - c\}$. In the subgame, neither goes to court for $\omega \leq x + c$ and Firm goes to court for $\omega > x + c$.

3. It is optimal for Firm to generate no information and induce mutual silence in the subgame if $\max\{c, 1-, c\} \leq x \leq \mathbb{E}[\omega] + c$.

4. It is optimal for Firm to learn ω above $x + c$ and no other information if $0 < x - c < \mathbb{E}[\omega|\omega \leq x + c] < \mathbb{E}[\omega]$. In the subgame, neither goes to court for $\omega \leq x + c$ and Firm goes to court for $\omega > x + c$. This is the low-risk case.

5. It is optimal for Firm to learn ω above s^* and no other information if $\mathbb{E}[\omega|\omega \leq x + c] \leq x - c \leq \mathbb{E}[\omega]$. In the subgame, neither goes to court for $\omega \leq s^*$ and Firm goes to court for $\omega > s^*$. This is the medium-risk case.

6. It is optimal for Firm to learn ω below s_* and no other information if $\mathbb{E}[\omega] < x - c$. In the subgame, neither goes to court for $\omega \geq s_*$ and Regulator goes to court for $\omega < s_*$. This is the high-risk case.

C Comparative Statics

In Proposition 5, I discussed the effect of status quo x on Firm's utility. In this section, I assess the effect of cost c on Firm's utility. In the low-risk equilibrium, Firm receives

$$\int_0^{x+c}(x)dF(\omega) + \int_{x+c}^{1}(\omega - c)dF(w).$$

Consider a larger $c' > c$ such that results in the low-risk equilibrium. Firm receives

$$\begin{aligned}
&\int_0^{x+c'}(x)dF(\omega) + \int_{x+c'}^{1}(\omega - c)dF(w) \\
=&\int_0^{x+c}(x)dF(\omega) = \int_0^{x+c'}(x)dF(\omega) + \int_{x+c'}^{1}(\omega - c)dF(\omega) \\
<&\int_0^{x+c}(x)dF(\omega) = \int_0^{x+c'}(\omega - c)dF(\omega) + \int_{x+c'}^{1}(\omega - c)dF(\omega) \\
=&\int_0^{x+c}(x)dF(\omega) + \int_{x+c}^{1}(\omega - c)dF(w),
\end{aligned}$$

which is smaller than the payoff induced by the original c. In other words, conditional on the low-risk equilibrium, Firm's utility *decreases* in c.

In the medium-risk equilibrium, Firm receives

$$\int_0^{1}(\omega)dF(\omega) + 2\int_0^{s^*}(c)dF(\omega) - c.$$

Note that s^* decreases in c because the right-hand side of condition $\mathbb{E}[\omega|\omega \leq s^*] = x - c$ decreases in c. Then, indirect and direct effects of c on the second term of the expression are countervailing. In other words, in the medium-risk equilibrium, the effect of c cannot be determined without further assuming the shape of prior distribution F.

In the high-risk equilibrium, Firm receives

$$\int_0^{1}(\omega)dF(\omega) + \int_{s_*}^{1}(c)dF(\omega).$$

Note that s_* decreases in c. Therefore, conditional on the high-risk equilibrium, Firm's utility increases in c.

D Exogeneous Absence of Information

In this section, I compare Regulator's utilities between Firm's choice of information and the exogenous absence of information. In the low-risk equilibrium, Regulator receives

$$\int_0^{x+c} (1-x)dF(\omega) + \int_{x+c}^1 (1-\omega)dF(\omega).$$

Suppose $\mathbb{E}[\omega] > x+c$. I have established that Regulator prefers Firm's choice of information if $\mathbb{E}[\omega|\omega < x+c] > x$ and no information otherwise. Suppose $\mathbb{E}[\omega] \leq x+c$. Then, neither goes to court without information. Regulator receives a payoff of $1-x$. In other words, Firm *does not* go to court for $\omega > x+c$ unlike in the baseline model. Regulator prefers no information for this reason.

In the medium-risk equilibrium, Regulator receives

$$\begin{aligned}
&\int_0^{s^*} (1-x)dF(\omega) + \int_{s^*}^1 (1-\omega)dF(\omega) \\
=&1 - \int_0^{s^*} (x)dF(\omega) - \int_{s^*}^1 (\omega)dF(\omega) \\
=&1 - \int_0^{s^*} (\omega+c)dF(\omega) - \int_{s^*}^1 (\omega)dF(\omega) \\
=&1 - \int_0^1 (\omega)dF(\omega) - \int_0^{s^*} (c)dF(\omega).
\end{aligned}$$

Suppose $\mathbb{E}[\omega] > s^*$. Then, Firm goes to court and Regulator receives a payoff of $1-\mathbb{E}[\omega]$. Regulator prefers no information. If $\mathbb{E}[\omega] \leq s^*$, neither goes to court without information. Regulator receives a payoff of $1-x$. Similar to the low-risk case, Regulator prefers no information.

In the high-risk equilibrium, Regulator receives

$$\begin{aligned}
&\int_0^{s_*}(1-\omega-c)dF(\omega)+\int_{s_*}^1(1-x)dF(\omega)\\
=&1-\int_0^{s_*}(x)dF(\omega+c)-\int_{s_*}^1(x)dF(\omega)\\
=&1-\int_0^{s_*}(x)dF(\omega+c)-\int_{s_*}^1(\omega+c)dF(\omega)\\
=&1-\int_0^1(\omega)dF(\omega)-c.
\end{aligned}$$

Without information, Regulator goes to court and receives a payoff of $1-\mathbb{E}[\omega]-c$. Regulator is indifferent between Firm's choice of information and the exogeneous absence of information.

E Example: Binary State

Consider the limit case of the benchmark model where state ω is binary. It is straightforward to apply the derived implications from the main analysis to this setting. However, it is useful to solve the model again with a concavification method as observed by Kamenica and Gentzkow (2011).

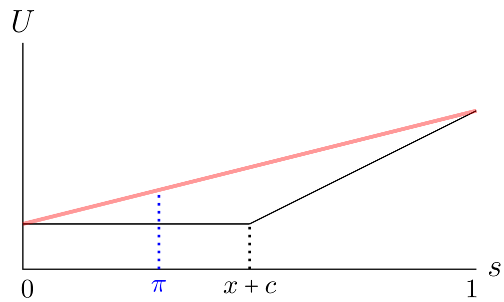

(a) $x - c < 0$ and $x + c < 1$

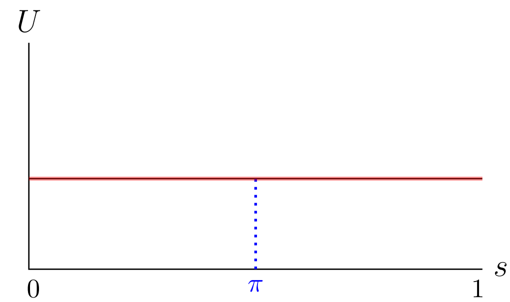

(b) $x - c < 0$ and $x + c > 1$

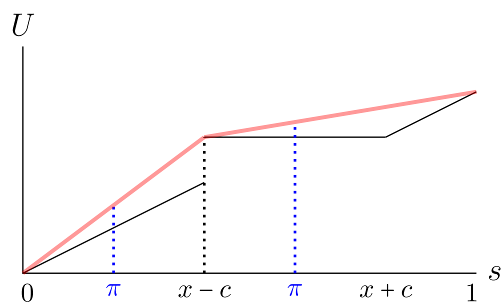

(c) $x - c > 0$ and $x + c < 1$

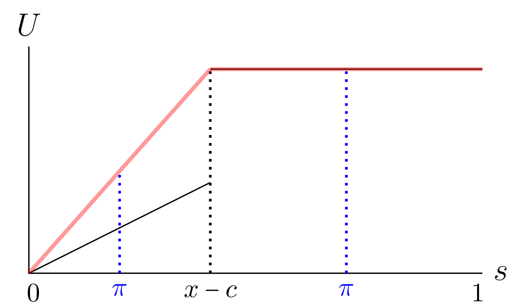

(d) $x - c > 0$ and $x + c > 1$

Figure E.4: These figures plot Firm's indirect utility function $U(s)$ (black sold) and the smallest concave function that uniformly stays above $U(s)$ (red translucent).

1. Firm's optimal choice of signal does not matter if $x \in (1 - c, c)$ because both players remain silent for all beliefs.

2. It is optimal for Firm to fully learn ω if $x \leq \min\{c, 1-c\}$. In the subgame, both players remain silent for $\omega = 0$ and Firm goes to court for $\omega = 1$.

3. It is optimal for Firm to provide no information and induce mutual silence in the subgame if $x \geq \max\{c, 1 - c\}$ and $\pi \geq x - c$.

4. It is optimal for Firm to partially learn $\omega = 1$ so that induced posterior belief (absent information) is $x - c$ if $x \in (c, 1 - c)$ and $\pi \geq x - c$. In the subgame, Firm goes to court

when $\omega = 1$ is revealed, and both players remain silent otherwise.

5. It is optimal for Firm to partially learn $\omega = 0$ so that induced posterior belief (absent information) is $x - c$ if $\pi < x - c$. In the subgame, Regulator goes to court when $\omega = 0$ is revealed, and both players remain silent otherwise.

Note that low-risk case does not exist in a binary setting. Recall that the Bernoulli distribution is an extreme case of dispersion given the mean; the risk is never low enough to satisfy the condition $\mathbb{E}[\omega|\omega \leq x + c] > x - c$. Based on this binary ω setting, I examine three extensions of the model. The first extension with the pre-game settlement is presented in Proposition 7. I present the other two below.

E.1 Privately Informed Firm

In this section, I present Firm's interim optimal choice of signal when ex-ante optimal choice unravels given private information. Recall that the ex-ante optimal experiment unravels for high-risk (partially reveal $\omega = 0$) and medium-risk (partially reveal $\omega = 1$) cases.

Suppose $\pi < x - c$ and $x + c < 1$. Then, Firm's ex-ante optimal choice is to reveal $\omega = 0$ with probability q_0 such that satisfies $\pi/(\pi + (1 - \pi)(1 - q_0)) = x - c$. In equilibrium, Regulator goes to court when $\omega = 0$ is revealed, and both remain silent otherwise. Therefore, Firm knowing $\omega = 0$ receives $(1 - q_0)x$ and Firm knowing $\omega = 1$ receives x in equilibrium. If Firm fully discloses, she instead receives 0 and $1 - c$ for each ω. Therefore, Firm knowing $\omega = 1$ has the incentive to deviate from the ex-ante optimal mechanism and reveal her knowledge. The signal can be interim optimal by choosing smaller $q < q_0$ because Firm can lower the posterior belief mean enough to *induce* Regulator to go to court by doing so. Then, Firm knowing $\omega = 0$ receives 0 and Firm knowing $\omega = 1$ receives 1. Essentially, it is interim optimal for Firm to not provide any information.

Now suppose $0 < x - c \leq \pi$ and $x + c < 1$. Then, Firm's ex-ante optimal choice is to reveal $\omega = 1$ with probability q_1 such that satisfies $\pi(1 - q_1)/(\pi(1 - q_1) + 1 - \pi) = x - c$. In equilibrium, Firm goes to court when $\omega = 1$ is revealed, and both remain silent otherwise. Therefore, Firm knowing $\omega = 0$ receives x and Firm knowing $\omega = 1$ receives $q_1(1-c)+(1-q_1)x$

in equilibrium. As before, Firm knowing $\omega = 1$ has the incentive to deviate from the ex-ante optimal mechanism and reveal her knowledge. Note that Firm receives x when ω is not revealed because Regulator chooses to remain silent when the posterior mean is $x - c$. However, Regulator is, in fact, indifferent between remaining silent and going to court here. By assuming that Regulator goes to court instead, the same strategy becomes interim optimal because Firm now receives $q_1(1-c) + 1 - q_1$. This is the ex-ante optimal interim optimal because q_1 maximizes the probability that Firm receives 1, not $1 - c$.

E.2 "Getting It Right"

Additionally, I consider an extension where Regulator's objective is to "get it right" rather than to pursue $\omega = 0$. Then, the payoff structure of the conflict stage is as below.

Firm \ Regulator	N	C
N	$(x, -\lvert x-\omega\rvert)$	$(\omega, -c)$
C	$(\omega - c, 0)$	$(\omega - c, -c)$

Table E.1: Payoff Structure

Nash equilibria for the conflict stage are as follows.[1]

1. Both remain silent if $(2x-1)\mathbb{E}[\omega|s] \geq x - c$.

2. Firm remains silent and Regulator goes to court if $(2x-1)\mathbb{E}[\omega|s] < x - c$.

The indirect utility of Firm is

$$U(s) = \begin{cases} x & \text{if } (2x-1)s \geq x - c \\ s & \text{if } (2x-1)s < x - c \end{cases}$$

As in the benchmark model, $U(s)$ can take largely four different shapes depending on x and c. When $x \in (1-c, c)$, $U(s)$ is constant at x. Therefore, Firm's choice of G does not matter

[1]Technically, another subgame equilibrium exists: Firm goes to court, and Regulator remains silent if $\mathbb{E}[\omega|s] > x + c$. This parameter space is a subset of the parameter space that induces equilibrium where Regulator goes to court instead. Therefore, I assume that Regulator goes to court when a multiplicity of equilibria exists.

here; any s leads to a subgame equilibrium where both remain silent. When $x \leq \min\{c, 1-c\}$, Firm's optimal choice of G is to fully reveal ω by connecting $(0, x)$ and $(1, 1)$. This strategy is the same as that from the benchmark model given same x and c. When $x \in (c, 1-c)$, the utility is linear and continuous. Therefore, Firm's choice of G does not matter; any s leads to a subgame equilibrium where Regulator goes to court.

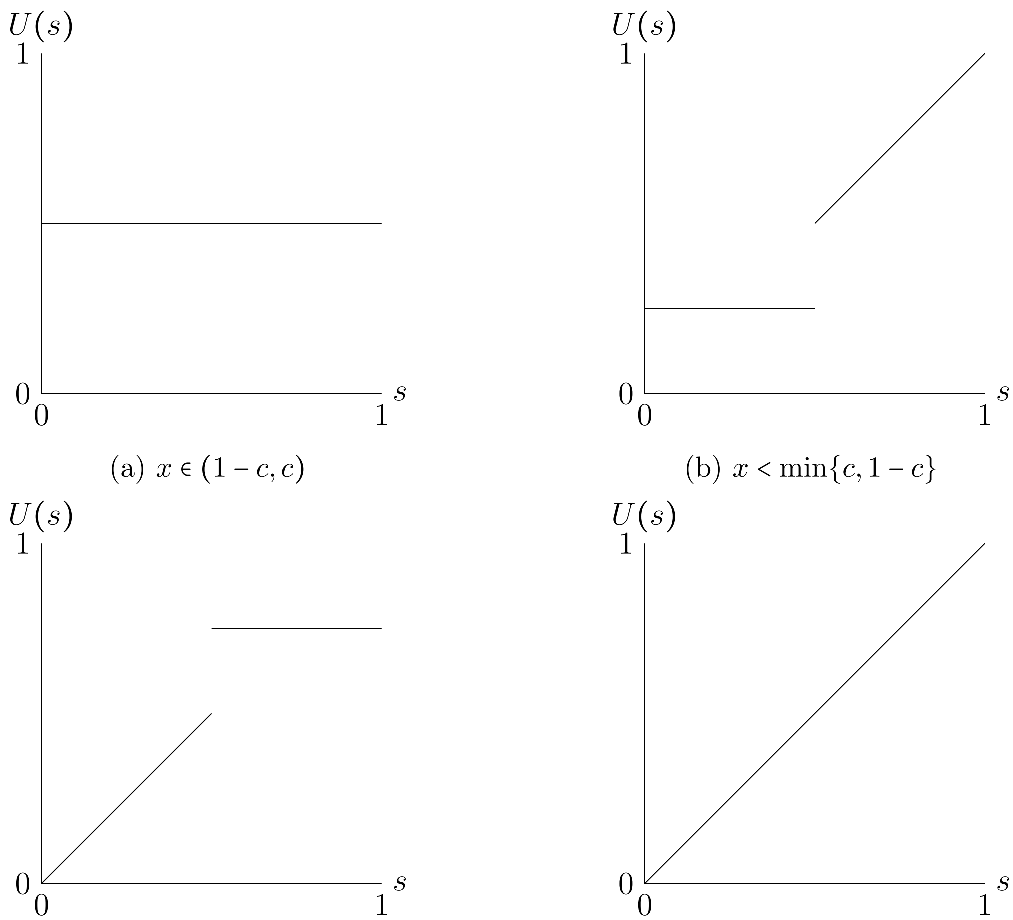

(a) $x \in (1-c, c)$

(b) $x < \min\{c, 1-c\}$

(c) $x > \max\{c, 1-c\}$

(d) $x \in (c, 1-c)$

Figure E.5: This figure plots indirect utility depending on x and c.

When $x \geq \max\{c, 1-c\}$, there are four subcases:

- **Case 1.** If $\pi < (x-c)/(2x-1)$ and $(x-c)/(2x-1) < x$, Firm partially learns $\omega = 0$. If ω is realized, Regulator goes to court. Otherwise, both players remain silent with posterior belief $(x-c)/(2x-1)$.

- **Case 2.** If $\pi \geq (x-c)/(2x-1)$ and $(x-c)/(2x-1) < x$, Firm generates no information.

In the subgame equilibrium, both players remain silent.

- **Case 3.** If $\pi < (x-c)/(2x-1)$ and $(x-c)/(2x-1) \geq x$, Firm generates no information. In the subgame equilibrium, Regulator goes to court.
- **Case 4.** If $\pi \geq (x-c)/(2x-1)$ and $(x-c)/(2x-1) \geq x$, Firm partially learns $\omega = 1$. If ω is realized, both players remain silent. Otherwise, Regulator goes to court with posterior belief $(x-c)/(2x-1)$.

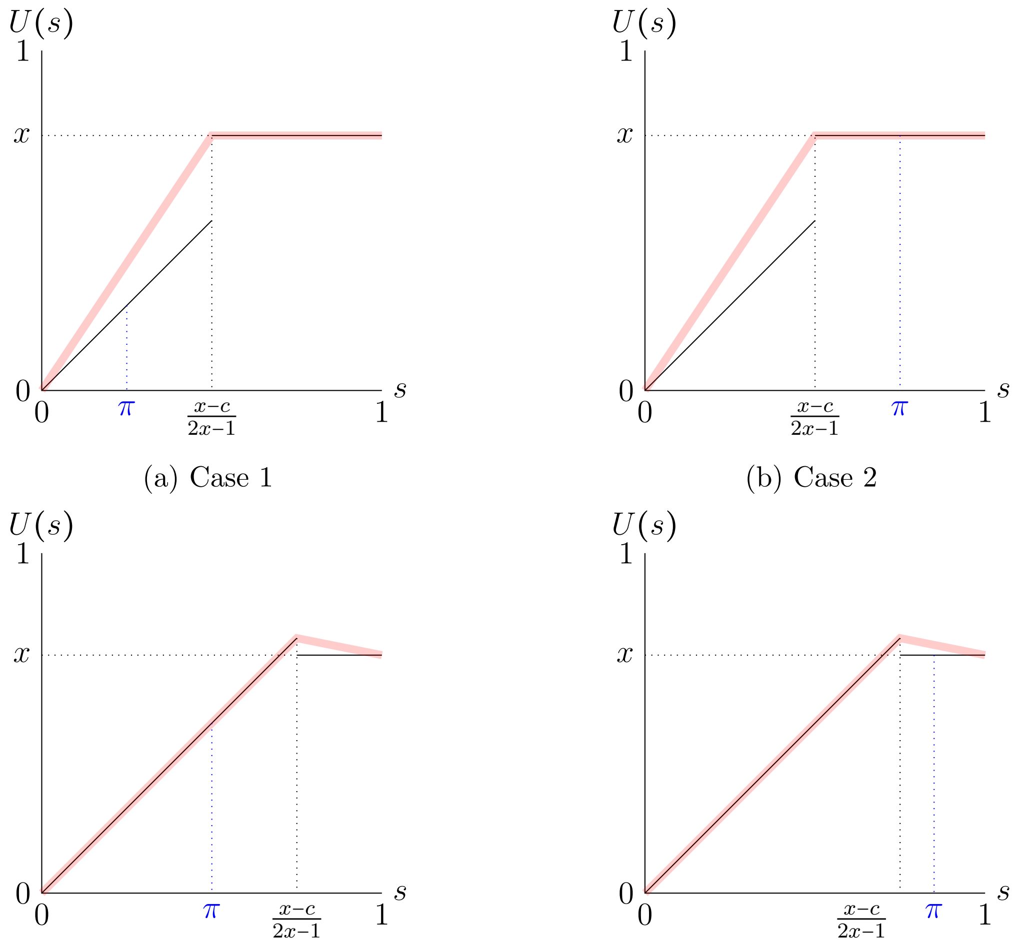

(a) Case 1 (b) Case 2

(c) Case 3 (d) Case 4

Figure E.6: This figure plots Firm's indirect utility given $x \geq \max\{c, 1-c\}$. The red translucent line is the smallest concave function that uniformly stays above $U(s)$.

F Heterogeneous Cost

I consider an extension where cost c is heterogeneous. Specifically, assume that $c_F = c$ and $c_R = c + \gamma$ where $c \in (0,1)$ and $\gamma \in (0, 1-c)$. This implies that $0 < c_F < c_R < 1$. Then, the payoff structure of the conflict stage is as described below.

Firm \ Regulator	N	C
N	$(x, 1-x)$	$(\omega, 1-\omega-c_R)$
C	$(\omega - c_F, 1-\omega)$	$(\omega - c_F, 1-\omega-c_R)$

Table F.2: Payoff Structure

Solving for the Nash equilibrium for the conflict stage,

1. Firm remains silent and Regulator goes to court if $\mathbb{E}[\omega|s] < x - c_R$.

2. Firm goes to court and Regulator remains silent if $\mathbb{E}[\omega|s] > x + c_F$.

3. Both remain silent if $\mathbb{E}[\omega|s] \in [x - c_R, x + c_F]$.

The corresponding indirect utility of Firm is

$$U(s) = \begin{cases} s & \text{if } s < x - c - \gamma \\ x & \text{if } s \in [x - c - \gamma, x + c] \\ s - c & \text{if } s > x + c. \end{cases}$$

Note that the only difference from the benchmark model is the threshold that deters Regulator from going to court. The equilibrium conditions are not too deviating either:

1. Firm's choice of signal does not matter if $x \in (1-c, c+\gamma)$ because both players remain silent for all beliefs.

2. It is optimal for Firm to fully reveal ω if $x \leq \min\{c+\gamma, 1-c\}$. In the subgame, both players remain silent for $\omega \leq x + c$ and Firm goes to court for $\omega > x + c$.

3. It is optimal for Firm to provide no information and induce mutual silence in the subgame if $\max\{c+\gamma, 1-c\} \leq x \leq \mathbb{E}[\omega] + c + \gamma$.

4. It is optimal for Firm to reveal ω above $x + c$ and provide no other information if $0 < x - c - \gamma < \mathbb{E}[\omega|\omega \leq x + c] < \mathbb{E}[\omega]$. In the subgame, both players remain silent for $\omega \leq x + c$ and Firm goes to court for $\omega > x + c$.

5. It is optimal for Firm to reveal ω above s^* and provide no other information, where $\mathbb{E}[\omega|\omega \leq s^*] = x - c - \gamma$. This is if $\mathbb{E}[\omega|\omega \leq x + c] \leq x - c - \gamma \leq \mathbb{E}[\omega]$. In the subgame, both players remain silent for $\omega \leq s^*$, and Firm goes to court for $\omega > s^*$.

6. It is optimal for Firm to reveal ω below s_* and provide no other information, where $\mathbb{E}[\omega|\omega \geq s_*] = x - c - \gamma$. This is if $\mathbb{E}[\omega] < x - c - \gamma$. In the subgame, both players remain silent for $\omega \geq s_*$, and Regulator goes to court for $\omega < s_*$.

As in the benchmark model, assume that $x \in (c + \gamma, 1 - c)$. It directly follows that as γ increases, the equilibrium changes to high-, medium-, and low-risk case.

Consider Firm's utility: Firm's utility in the low-risk case does not depend on γ. Firm's utility in the medium-risk case is $\mathbb{E}[\omega] + \mathbb{P}(\omega \leq s^*)(c + \gamma) + \mathbb{P}(\omega > s^*)(-c)$, which increases in γ because the probability of going to court increases in γ. Firm's utility in the high-risk case is $\mathbb{E}[\omega] + \mathbb{P}(\omega \geq s_*)(c + \gamma)$, which increases in γ. Next, consider Regulator's utility: Regulator's utility in the low-risk case does not depend on γ. Regulator's utility in the medium-risk case is $1 - \mathbb{E}[\omega] - \mathbb{P}(\omega \leq s^*)(c + \gamma)$, which decreases in γ. Regulator's utility in the high-risk case is $1 - \mathbb{E}[\omega] - c - \gamma$, which decreases in γ.