Judged in Hindsight: Regulatory Incentives in Approving Innovations

Suraj Malladi*

March 31, 2020

Abstract

I study how limited information and ex-post evaluation by third parties with the benefit of hindsight affect how regulators approve innovations. In the face of ambiguity over innovation characteristics, such a regulator limits or delays product approval, even when she is not waiting for new information to arrive. When evidence is costly for firms to generate but can be selectively reported, the regulator delegates information acquisition to the firm with the objective of minimizing max-regret. This model can explain observed patterns of correlation between firm costs and benefits of approval, why regulators drag their feet on approval decisions even in the face of strong favorable evidence, and support for regulatory sandboxes even when they do not hasten learning.

*Stanford Graduate School of Business. Email: surajm@stanford.edu. I thank my advisor, Andrzej Skrzypacz, and committee members Matthew Jackson and Michael Ostrovsky for their guidance at every step of this project. I am grateful for many rounds of feedback and suggestions from committee members Mohammad Akbarpour, Kyle Bagwell and Yonatan Gur. I also thank Jeremy Bulow, Daniel Chen, Zi Yang Kang, Edward Lazear, Yucheng Liang, Andres Perlroth, Marcos Salgado, David Yang, Political Economy Theory Seminar participants at Stanford GSB and especially Anirudha Balasubramanian for many helpful discussions.

1 Introduction

Regulatory agencies charged with ensuring safety or efficacy of innovations are often ill-informed and ill-equipped for the task. They know little about product characteristics, and much of their information must come from parties with incentives to disclose only favorable evidence. Despite the difficulties of consumer protection, various media, consumer watchdogs and other political actors tend to measure the actions of regulators against the best approval decision in hindsight. Congress may investigate the FDA for failing to recall unsafe drugs in a timely manner, while patient advocacy groups hold it accountable for delaying the passage of life-saving drugs. The National Highway Traffic Safety Administration (NHTSA) may be criticized for delaying approval of automated vehicles and perhaps criticized again if the technology turns out to cause an uptick in accidents.

Facing ambiguity, cherry-picked evidence, and ex-post scrutiny, how does a regulator design an approval process for innovations? How does design change with the metric of ex-post evaluations, i.e., consumer welfare or consumer welfare and regulatory costs?

In my model, a regulator commits to an approval rule for new innovations, which maps the evidence a firm shows to its level of approval. A firm generates costly evidence by running trials. Each trial results in a success or failure with frequencies depending on its privately known innovation quality. The firm reports a favorable subset of this evidence to get some level of market approval. The quality of an innovation is eventually revealed regardless of the regulator's decision. The regulator is then evaluated on the basis of her approval decision and ex-post optimal action (immediate approval or rejection) given the status quo outcome and innovation quality. Her anticipation of potential criticism, given the yardstick by which her performance is measured, guides her choices.

The regulator only knows the space of possible firm types but does not have a prior on this space. She instead takes actions to maximize her worst-case performance on ex-post evaluations. The firm knows its chances and costs of producing successful trials. It has some prior over its true quality (and therefore its potential benefits), but the experiments it runs for the regulator give no new information to the firm.[1]

Consider first the case where the firm can costlessly fabricate evidence of any sort, so no information can be credibly communicated. Suppose that evaluations punish the regulator by the size of her error: how much worse (better) is the approved (rejected) innovation than the status quo. If the regulator has to make immediate or full approve or reject decisions, she is limited in her ability to cope with ambiguity. If she passes a drug, her worst outcome is that it turns out to be very harmful or ineffective, and vice versa if she rejects. Depending on which error is punished less, she either approves all

[1]The FDA, for example, may ask a drug company to show efficacy of a heart disease treatment in terms of its performance on a surrogate measure like cholesterol-reduction. The drug company may already know how effectively it performs on this short term metric, and running trials may not reveal how effectively the drug treats heart disease.

innovations or waits for the eventual revelation of quality before passing any innovation. These regulatory modes are sometimes referred to as 'permissionless innovation' or 'strict precautionary principle', respectively.

But regulators can often partially approve innovations (e.g., through production quotas, allowing sales in only some markets or delaying approval), and ex-post criticisms may scale with the extent of approval. Consider for example an FDA regulator who is judged, at the time the public learns the drug's quality, by how long she had delayed approval of a good drug or left a bad drug on the market. [2] Such a regulator can mitigate the size of her worst case error by awarding partial approval. A recent political theory literature emphasizes reputation to multiple audiences as the main driver of behavior among government bureaucracies (Carpenter, 2004). These incentives coupled with ambiguity over innovation quality can explain why regulators limit or delay approval even when they are not collecting or waiting for information.

When evidence is costly to generate and successful trials cannot be fabricated, the regulator can reduce her worst case error further by committing to rules which vary the approval awarded by the evidence reported. Even when firms disclose one-sided evidence, costs of generating positive evidence are lower and benefits of approval higher for firms with better quality innovations. Muddling this relationship is the fact that some firms may simply enjoy lower costs of running trials or better marketing of their products. Under any approval rule, firms with the same ratio of costs to benefits (i.e., costs of running successful trials to benefits of approval) will choose to provide the same levels of evidence to the regulator. Therefore, screening in this multidimensional type space reduces down to these one-dimensional ratios, which I call *effective types*. This is consistent with evidence from FDA trials that approval delays (occurring *after* firms complete and submit all trial results) are positively correlated to the ratio of costs to revenues but insignificantly correlated to costs alone (Olson, 1997).

When the regulator is evaluated on the basis of consumer protection, the worst case error of passing firms with low benefit to cost ratios can be smaller than the worst case error of passing firms with intermediate benefit to cost ratios. For example, a firm may have low benefit of approval for one of two reasons: its innovation is of obviously low quality and would have low demand, or its innovation is good but has a niche market. Approval is good in the latter case, and the harm of approving is limited in the former case. On the other hand, a firm with an intermediate benefit of adoption may have an innovation that is worse than the status quo, but not noticeably so. If such a product enjoys high demand, it can be more destructive to consumer welfare. In such cases, a regulator commits to giving *less approval* to low effective types than she would have given had effective types been observable.

In some instances, the costs of approval incurred by regulated firms are also revealed,

[2] Or alternatively, consider a regulator who is evaluated by what portion of markets she exposes to a bad AV technology or withholds a good one.

and ex-post evaluations may punish regulators for such costs. For example, patient groups worried about the spill-over into drug prices often criticize the drug approval process in cases where clinical trial costs are revealed to be large. When evaluations punish regulatory costs and errors in consumer protection, a regulator commits to awarding *more approval* to low effective types than she would have given had effective types been observable. On one hand, she risks making larger type I errors on low effective types. But the benefit is that this reduces the burden of evidence for effective types to whom the regulator wishes to give more approval, i.e., for whom the potential error of rejection is worse than the error of acceptance.

In fact, the regulator optimally awards some approval to a firm that shows no evidence, which may explain widespread bureaucratic support for approval rules involving regulatory sandboxes (see, for example, Buckley et al. (2019)). Next, the option of full approval may not exist in a regulator's optimal evidence-approval menu, even if firms are willing to show evidence that guarantees their superiority to the status quo. This happens when the regulator's loss from allowing high effective types to incur more costs and separate themselves further outweighs the benefit of making smaller type II errors. It behooves the regulator to partially approve and wait for the eventual revelation of the firm's true quality before revising her earlier approval decision.

The next subsection reviews existing literature. Section 2 solves a simplified version of the model to show how ex-post evaluations can lead to partial approval or delays. Section 3 describes the full model, which can explain observed variation in approval times. Section 4 solves for optimal rules of a regulator who cares only about consumer protection. Section 5 characterizes regulator optimal approval rule when the regulator is also penalized for costs imposed on innovating firms. Section 6 considers how the forms of optimal approval rules compare when alternative assumptions about the regulator's information or objectives are made. It also discuss why the agency and communication problems modeled here are difficult to overcome in practice. Section 7 concludes.

1.1 Related Literature

Like McClellan (2017), Escobar et al. (2019) and Tetenov (2016), the focus of this paper is to characterize optimal approval rules for a regulator who has commitment power, facing an innovating firm that can undertake costly experimentation. McClellan (2017) considers a setting where firms decide when to stop public experimentation, and regulators lower standards of approval over time to encourage continued experimentation.[3] Escobar et al. (2019) study a model where firms with private information experiment privately and communicate with a regulator via cheap talk. To incentivize truthful reporting, the regulator commits to rules that delay approval even after it becomes common knowledge

[3] Henry and Ottaviani (2019) compare various arrangements of commitment power between a regulator and firm in a similar model. Unlike McClellan (2017), they study a setting with common knowledge and restrict to stationary approval rules in the case with regulator commitment.

that a project is good. Regulators in my setting delay or partially approve both to provide incentives *and* to mitigate worst-case criticism. Closest to my setting is the static model in Tetenov (2016), who also studies how regulators facing ambiguity design approval rules when firms can provide verifiable and costly evidence. In my model, multidimensional types and a concern for minimizing regret preclude the possibility of perfectly screening out undesirable innovations, as in Tetenov (2016).[4]

The current paper departs from earlier work on optimal approval rules in two substantive ways. First, I assume that firms can manipulate by hiding negative trials or manipulating experiments to increase chances of successful trials. These modes of manipulation are closest to the models of Janssen (2019) and Di Tillio et al. (2017), both of which study the regulator's approval decision without commitment. Second, I assume that regulators care about ex-post evaluations, which may take into account correctness of decisions and costs imposed on firms. Unlike the aforementioned papers, both the regulator and the firm may incur costs when the latter generates evidence.[5]

The regulator's concern for criticism in my model is motivated by a growing literature on bureaucratic reputation, notably Olson (1995), Carpenter (2014), Moynihan (2012), and Gilad et al. (2015). These papers show evidence that reputation before multiple audiences drives regulators to behave in a way so as to avoid negative feedback.[6] Additionally, Carpenter (2004) argues that some of these audiences have increased the visibility (and therefore punishment) of type II errors. This contrasts with the older view, famously argued by Milton Friedman, that regulators worry only about mistaken approvals (Friedman, 1973) and motivates the objective function highlighted here.[7]

The solution to regulator's problem, under ambiguity, is consistent with some observed evidence on approval processes. Carpenter (2002) shows that significant delays in FDA review occur *after* clinical trials are completed, though nearly all drugs that reached this stage eventually get approved. Olson (1997) shows that the substantial variation in FDA delay times is correlated to a ratio of firm research costs to post-approval revenue, controlling for many other firm and drug characteristics, but uncorrelated to costs alone. Finally, Buckley et al. (2019) documents how FinTech regulators in many countries give away limited approval without much prior evidence in regulatory sandboxes. Sometimes regulators push for these measures even when there is little learning in practice. The model studied here accounts for all of these patterns. The same model with a Bayesian regulator or a regulator concerned directly with product quality would not.

[4]Information generated by the firm is muddled as in Frankel and Kartik (2019). Unlike their setting, the regulator has commitment power, so the focus here is on screening firms rather than signaling.

[5]Compare this to Guo (2016), for example, in which only the principal delegates experimentation but also solely incurs experimentation costs.

[6]This contrasts with theories that focus on regulatory capture, budget maximization, or treat regulators as social planners. See Noll (1982).

[7]This objective is reminiscent to that of Banerjee et al. (2017), which studies how researchers choose samples to convince audience members with worst-case beliefs. But the current paper is about regulators who, unlike researchers, delegate experimentation and contend with manipulation.

Like Carroll (2015) and Frankel (2014), I use ambiguity to derive features of real world contracts. Within the robust mechanism design literature, the current paper is especially related to those with min-max regret criteria like Bergemann and Schlag (2008) and Guo and Shmaya (2019), who respectively study robust monopoly pricing and monopoly regulation. In comparison, I study consumer protection and design of approval processes.

2 Preliminary Model Without Verifiable Evidence

Before moving to the full model, it is instructive to consider a simpler model in which no information is communicated between the regulator and the firm. This helps highlight the role of ex-post evaluations on the approval decisions of the regulator.

The firm has an innovation of type $\theta \in \Theta$ and seeks approval from a regulator to market its product. We assume that any mappings that we later define on Θ have finite maxima and minima and expectation, whenever we invoke them. The firm receives a payoff of $b(\theta) > 0$ if its innovation is approved by the regulator, and it receives 0 otherwise.

The regulator chooses an 'extent of approval', $y \in [0, 1]$. One interpretation of y is that it is the proportion of markets in which the firm is allowed to sell its product. $y = 1$ means the firm has full approval, and $y = 0$ means the firm is banned from selling its innovation anywhere. $y \in (0, 1)$ means the firm restricting output or selling in some markets but not others, e.g., the Drug Enforcement Agency regulator sets opioid production quotas or the NHTSA allows automated vehicles to be driven in some districts but not others.

In other contexts, the extent of approval can be thought of as the level of delay a firm has to face prior to having a product approved. $y = 1$ means a drug is approved immediately, $y = 0$ means a drug is never approved, and $y \in (0, 1)$ means a drug is approved after a period of delay. For example, there can be a flow of patients suffering from a particular disease who need immediate treatment. A drug that is immediately approved is accessible to patients immediately. If the drug is approved only with delay, it is effectively denied to the patients who would have used it in the intervening time. Delay also diminishes the profits of the pharmaceutical company which paid research and development costs upfront but can only begin to enjoy sales at a future date.

In this section, we assume that the firm's scope for manipulating the evidence is very large. In fact, the only way the firm can communicate with the regulator is through cheap talk. By our assumptions, the firm always prefers approval to no approval, so only a babbling equilibrium exists: no information is conveyed between the firm and the regulator.

The regulator has a payoff that depends on the type of the firm's product and scales with the extent of approval. Let $\tilde{v}(\theta) \in \mathbb{R}$ be some aggregate measure of the innovation's quality (which could account for its safety, efficacy, etc.) and let $\tilde{v}_0$ be the quality of the existing outside option. Before turning to the case where the regulator is worried about

ex-post evaluations, first consider the case of a regulator who is directly concerned with ensuring a larger extent of the market utilizes a higher quality product. In particular, a regulator who chooses an approval level y for an innovation of type θ receives a payoff of

$$y\tilde{v}(\theta) + (1-y)\tilde{v}_0.$$

The regulator wants to approve the innovation only when $\tilde{v}(\theta) \geq \tilde{v}_0$. But while all other aspects of the game are common knowledge (i.e., Θ, $\tilde{v}_0$, and the mappings b and $\tilde{v}$), θ is privately known to the firm and not to the regulator.

2.1 Average or Worst-Case Product Quality

Suppose first that the regulator has some prior F over Θ. Then the regulator will award full approval if $E_F[\tilde{v}(\theta)] > \tilde{v}_0$, reject the innovation if $E_F[\tilde{v}(\theta)] < \tilde{v}_0$, and is indifferent to any level of approval when $E_F[\tilde{v}(\theta)] = \tilde{v}_0$.

Taking one step closer to the setting considered in this paper, suppose instead that the regulator faces ambiguity over Θ and evaluates any decision by its worst case outcome.[8] Such a regulator will approve only if there is no possibility that the quality of the innovation is worse than the status quo, i.e., $\inf_{\theta\in\Theta} \tilde{v}(\theta) \geq \tilde{v}_0$. As in the Bayesian case, the optimal policy can be implemented as an immediate rejection or approval.

2.2 Ex-post Evaluations

Now consider instead the case where the regulator is concerned, not with product quality directly, but by how she is evaluated by third parties in hindsight. Carpenter (2004) emphasizes the primacy of such concerns in driving incentives:

> Bureaucratic reputations often have multiple audiences... Reputation is simply a currency of bureaucratic politics. Agencies with strong reputations can more easily attract desired personnel; fend off budget cuts; and lobby for the programs, funds, and other things they desire. There are other things that bureaucracies protect and "maximize," but for many agencies such as the FDA, reputation protection serves as the simplest and most powerful dynamic governing their behavior.

Suppose that the state of the firm's innovation is exogenously revealed to the regulator and third-parties at some point after it comes to the regulator for review. In particular, suppose that the regulator's chosen extent of approval does not affect the timing of

[8] There are many examples where regulators either claimed or were directed to behave according to a maximin principle. For example, a 1981 White House directive explicitly directed all federal agencies to declare and evaluate decisions by "reasonable projections of the worst possible consequences of a proposed action" in issues such as environmental regulation or storage plans for nuclear waste, where "catastrophic uncertainty" was involved (Farber, 2010). Cass Sunstein documents the ubiquity of this sort of decision regulatory making under ambiguity in his book *Worst-Case Scenarios* (Sunstein, 2009).

the revelation of the state. This assumption embodies two common features of many regulatory approval processes.

The first feature is weak post-market oversight: consequences of approval are difficult to actively monitor (see Section 6). As in the recent public awakening to a decades long opioid crisis, the quality of a product is publicly understood and the relevant agency undertakes costly hearings and clean-up measures long after the initial approval decision.

The second feature is that exogenous revelation of information is possible even when a regulator rejects an innovation. For example, a drug or a technology that is blocked in one jurisdiction may still be approved in other countries and its qualities may become publicly revealed. Additionally, Carpenter (2004) argues that in the FDA context, the "rise in patient advocacy has led to a balancing of the visibility of Type II versus Type I errors", suggesting other parties could eventually produce unbiased information even when the firm does not.

When the state of the world, θ is revealed, the regulator is evaluated on how large an error she had made in her approval decision. Her error scales with her extent of approval. The regulator faces faces ambiguity over Θ and attempts to minimize worst case payoff over this space, i.e., solves

$$\inf_{y\in[0,1]} \sup_{\theta\in\Theta} f(\tilde{v}_0 - \tilde{v}(\theta))y\mathbb{1}_{\tilde{v}_0\geq\tilde{v}(\theta)} + g(\tilde{v}(\theta) - \tilde{v}_0)(1-y)\mathbb{1}_{\tilde{v}_0<\tilde{v}(\theta)},$$

where $f, g : \mathbb{R}_+ \to \mathbb{R}_+$ are strictly increasing with $f(0) = g(0) = 0$. These loss functions capture how severely errors of each kind are punished and may even encode the probabilities with which errors are observed, depending on their size.

The objective of minimizing worst case type I and type II errors is equivalent to minimizing *max regret*, after a change of variables. To see this, define $v : \Theta \to \mathbb{R}$ such that $v(\theta) \equiv (\tilde{v}_0 + C) - f(\tilde{v}_0 - \tilde{v}(\theta))$ if $\tilde{v}(\theta) \leq \tilde{v}_0$ and $v(\theta) \equiv g(\tilde{v}(\theta) - \tilde{v}_0) + (\tilde{v}_0 + C)$ if $\tilde{v}(\theta) > \tilde{v}_0$. Without loss of generality, suppose $C > 0$ was chosen such that $v \geq 0$; denote $\tilde{v}_0 + C$ by v_0. A regulator who minimizes worst case losses from type I and type II errors chooses $y \in [0,1]$ to solve

$$\begin{aligned}
&\inf_{y\in[0,1]} \sup_{\theta\in\Theta} f(\tilde{v}_0 - \tilde{v}(\theta))y\mathbb{1}_{\tilde{v}_0\geq\tilde{v}(\theta)} + g(\tilde{v}(\theta) - \tilde{v}_0)(1-y)\mathbb{1}_{\tilde{v}_0<\tilde{v}(\theta)} \\
&= \inf_{y\in[0,1]} \sup_{\theta\in\Theta} |v_0 - v(\theta)|y\mathbb{1}_{v_0\geq v(\theta)} + |v_0 - v(\theta)|(1-y)\mathbb{1}_{v_0<v(\theta)} \\
&= \inf_{y\in[0,1]} \sup_{\theta\in\Theta} \max\{v_0, v(\theta)\} - (1-y)v_0 - yv(\theta) \\
&\equiv \inf_{y\in[0,1]} R(y)
\end{aligned}$$

For any extent of approval and type θ, the regulator's utility is the difference between the 'payoff' to the best action in hindsight and the 'payoff' to the action taken by the regulator, where this 'payoff' is measured by v. Since the regulator evaluates decisions

by worst case, she imagines an adversary who chooses θ to maximize her regret, and we denote this max regret by $R(y)$.[9]

To give the regulator's optimal decision rule, we define $l \equiv \inf v(\Theta)$ and $m \equiv \sup v(\Theta)$.

LEMMA **1.** *Let y^* be a solution to $\min_{y\in[0,1]} R(y)$, and suppose $m > l$.*

1. *If $l > v_0$, $y^* = 1$.*
2. *If $m < v_0$, $y^* = 0$.*
3. *Otherwise, $y^* = \frac{m-v_0}{m-l}$.*

Proof. If $l > v_0$, $R(y) = \sup_{\theta\in\Theta}(1-y)(v(\theta)-v_0)$, which is minimized at $y^* = 1$. Case 2 is handled symmetrically.

For case 3, note that

$$\begin{aligned} R(y) &= \sup_{\theta\in\Theta} \max\{y(v_0 - v(\theta)), (1-y)(v(\theta)-v_0)\} \\ &= \max\{\inf_{\theta\in\Theta} y(v_0 - v(\theta)), \sup_{\theta\in\Theta}(1-y)(v(\theta)-v_0)\} \\ &= \max\{y(v_0 - l), (1-y)(m-v_0)\}, \end{aligned}$$

which is minimized at y^*, where $y^*(v_0 - l) = (1-y^*)(m-v_0)$. □

Lemma 1 shows that a regulator who faces ambiguity and is worried about ex-post evaluation may find it favorable to partially approve innovations or only award approval after some delay.

2.3 Approval Delays in Practice

The preceding discussion suggests that responding to incentives arising from ex-post evaluations could be a reason for why regulators restrict approval or drag their feet with approval decisions. In many situations, natural competing explanations are that the regulator waits for information to arrive before taking an irreversible action or approves a product in a limited set of markets for the sake of experimentation. Henry Miller, a former FDA scientist, describes an instance where apparently concern for ex-post evaluations motivated one regulator to try and delay approval of a new type of insulin:

> My team and I were ready to recommend approval after four months' review. But when I took the packet to my supervisor, he said, "Four months? No way! If anything goes wrong with this product down the road, people will say we rushed it, and we'll be toast." That's the bureaucratic mind-set... A large part of regulators' self-interest lies in staying out of trouble. One way to do that, my supervisor understood, is not to approve in record time products that might experience unanticipated problems. (Miller, 2018)

[9] Appendix shows that a similar representation holds when $v(\theta)$ is stochastic.

In fact, motives for additional information acquisition alone seem insufficient in explaining systematic delays in the FDA review process, which occur *after* all evidence is submitted. Pharmaceutical companies that have made it to this stage almost always receive a decision of approval from the FDA. Notably, there is substantial variation in the length of the FDA review process ranging from a week to eighty two months after a company has submitted all evidence. Indeed, Carpenter (2002) writes:

> FDA review occurs only after pharmaceutical companies have completed three stages of clinical trials on the drug. The requirement that firms provide all information to the FDA drastically reduces the agency's information costs and reduces the variance of product uncertainty in drug review. Agency design, administrative procedures, or informational barriers will, by themselves, have difficulty explaining the variance in FDA review times.

3 The Model

With only cheap talk, there is no scope for the regulator to discriminate between firms. In reality, regulators typically require firms to generate evidence by running costly trials to signal their type, even if they can manipulate results by selectively disclosing evidence. Hereafter, we focus on such a model to explore how an ambiguity averse regulator incorporates the information disclosed by firms.

Olson (1997) shows that (after controlling for a host of other firm and drug characteristics) FDA review times tend to be negatively correlated with a *ratio* of research expenditure to aggregate drug sales, while the relationship between delays and aggregate research costs is statistically insignificant. The model here accounts for both of these patterns.[10]

3.1 Primitives, Firm Information and Firm Objectives

Players There is a regulator (she) and a firm (it) with an innovation.

The firm has a type $\theta \in \Theta$. The firm's type determines its marginal cost $k(\theta)$ of running trials, the probability $q(\theta)$ that each trial succeeds, and its private benefit $b(\theta)$ of approval. Moreover, the regulator has a payoff $v(\theta)$ to approving the firm. If the regulator does not approve the firm, her payoff is v_0 and the firm's payoff is 0. Let Z be the image of Θ under (k, q, b, v). I assume Z is a compact and connected set, bounded away from 0.

The firm's type is its private information. Θ, Z, k, q, b, v and v_0 are common knowledge. I defer describing the regulator's information about θ to Subsection 3.3.

[10] While Olson (1997) does not have data on the actual costs of clinical trials, she notes that "research-intensive firms may submit more comprehensive clinical evidence in support of [a New Drug Application]".

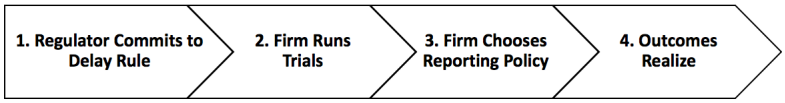

Figure 1: The game proceeds in four periods, with the regulator committing to a delay policy first.

Actions and Timing The game consists of four periods, as pictured in Figure 1. I describe these periods out of order, as it is convenient to first introduce the firm's actions.

In the second period, the firm runs trials in sequence and decides when to stop. Each trial has a binary outcome, resulting in a *success* with probability $q(\theta)$ or a *failure* with probability $1 - q(\theta)$.

After r rounds of trials, the firm privately observes the number of failures ρ_0^r and successes ρ_1^r that it has accumulated so far and chooses to run $\omega_r \in [0, 1]$ more trials.[11] Let $s_r \equiv (\rho_0^r, \rho_1^r)$ denote the firms cumulative results after r rounds of trials. If the firm chooses $\omega_r = 0$, the game moves to the third period. If $\omega_r > 0$, then the trial contributes to the cumulative outcomes, resulting in $s_{r+1} = (\rho_0 + \omega_r, \rho_1)$ with probability $1 - q(\theta)$ and $s_{r+1} = (\rho_0, \rho_1 + \omega_r)$ with probability $q(\theta)$.

Suppose the firm decides to stop running additional trials after r^* rounds, i.e., $\omega_{r^*} \in [0, 1]$. In the third period, the firm reports any tuple $(\hat{\rho}_0, \hat{\rho}_1) \in \mathbb{R}_+^2$ such that $\hat{\rho}_1 \leq \rho_1^{r^*}$: the firm can hide and fabricate failures but cannot fabricate successes. This describes the FDA drug approval process prior to the recent era where preregistration and transparency policies are seriously enforced; it can also capture the current state of AV regulation, for which formal safeguards have not yet been erected.[12]

In the fourth period, the firm's type θ is revealed and payoffs are realized. The revelation of the state happens independent of the level of manipulated evidence that the firm provides. One interpretation is that regardless of whether the drug was passed domestically, it would have been approved in a foreign country with potentially lower standards and the quality of the drug would be revealed. At this point, interest groups and media sources can evaluate the performance of the regulator. They can criticize the regulator for letting a bad drug linger on the market for too long or affect too many people, or criticize delayed approval of what turned out to be a good drug.

In the first period, the regulator commits to an *allocation rule* $x : \mathbb{R}_+^2 \to [0, 1]$, which is a mapping from the firm's reported successes and failures to an extent of market approval. In different cases, the regulator may modulate the extent of market approval by delaying

[11] Allowing for infinitely divisible trials makes results easier to state, and having discrete trials would not substantively affect results.

[12] Even if policies preclude burying negative evidence outright, a firm may still be able to inflate its success rate by choosing samples selectively or ending trials early at first indications of success. Appendix C shows how the model and results can be extended to treat such cases.

approval, opening an innovation to only a portion of the market, restricting production etc. I discuss these interpretations at the end of this section.

Firm's Objective A firm of type θ that runs $\rho_0^T + \rho_1^T$ trials and reports $(\hat{\rho}_0, \hat{\rho}_1)$ receives a payoff of

$$b(\theta)x(\hat{\rho}_0, \hat{\rho}_1) - k(\theta)(\rho_0^T + \rho_1^T),$$

and I assume the firm is risk neutral. Since the firm is free to report any $\hat{\rho}_0$, it produces evidence in stage one only to reach a target number of successes. Since evidence production is costly, it will produce exactly as many successes as it plans to at the outset. Note that the expected cost of producing one success is given by $c(\theta) \equiv \frac{k(\theta)}{q(\theta)}$.

Putting this together, I can express the firm's maximization problem as follows. A firm of type θ produces n successes to maximize its ex-ante expected payoff, i.e., to solve:

$$\sup_{n, \hat{\rho}_0 \in \mathbb{R}_+} b(\theta)x(\hat{\rho}_0, n) - c(\theta)n.$$

Before turning to the regulator's objectives, I show how the search for optimal approval rules can be simplified in my setting.

3.2 Simplifying the Space of Approval Rules

An allocation rule x induces a mapping from the firm's type to a corresponding optimal choice of extent of market approval and number of successful trials. The number of successful trials that a firm needs to show to achieve some level of market approval is a *standard of proof*. I let $\tilde{x}(\theta)$ and $\tilde{n}(\theta)$ denote (a selection of) a type θ firm's optimal choices given the regulator's allocation rule.

Let $u(\theta, \hat{\theta}) \equiv b(\theta)\tilde{x}(\hat{\theta}) - c(\theta)\tilde{n}(\hat{\theta})$. The regulator's allocation rule induces a menu consisting of tuples of market approval and standards of proof $(\tilde{x}, \tilde{n}) : \Theta \to [0,1] \times \mathbb{R}_+$ that satisfies incentive compatibility and individual rationality: for all $\theta, \hat{\theta} \in \Theta$,

$$u(\theta, \theta) \geq u(\theta, \hat{\theta}) \tag{IC}$$

$$u(\theta, \theta) \geq 0 \tag{IR}$$

An allocation rule $\tilde{x}$ is *implementable* if there exists some $\tilde{n}$ such that $(\tilde{x}, \tilde{n})$ is incentive compatible.

Let $\tau(\theta) \equiv \frac{b(\theta)}{c(\theta)}$, which is well defined since Z is bounded away from 0. Note that $b(\cdot)$ and $c(\cdot)$ are not co-monotone in general, the example in Section 4.1 illustrates.

LEMMA **2.** *If $\tilde{x}$ is implementable, and $\tau(\theta) > \tau(\hat{\theta})$, then $\tilde{x}(\theta) \geq \tilde{x}(\hat{\theta})$.*

This and other omitted proofs appear in Appendix E. The proof shows that all firm types θ with the same $\tau(\theta)$ have identical preferences over the menu of allocation probabilities and standards of proof offered.

Let $\mathcal{T} = \{\tau | \tau = \frac{b(\theta)}{c(\theta)} \text{ for } \theta \in \Theta\}$ be the space of the firm's *effective types*, which determine its behavior. Note that since $\mathcal{T}$ is a continuous transformation of Z, which is compact and connected, $\mathcal{T} = [\underline{\tau}, \overline{\tau}]$, for some $\underline{\tau}, \overline{\tau} \in \mathbb{R}_+$.

Assumption **1.** *All types of the firm break indifferences among optimal menu choices in the same way.*

This assumption limits the space of implementable allocation rules, but is not needed in any characterizations of optimal rules. With this, we can consider allocation rules and standards of proof as functions of effective types: $(y, n) : \mathcal{T} \to [0, 1] \times \mathbb{R}_+$. To avoid confusion with the term for mappings from Θ, I call such mappings *approval rules*. This leads to the familiar characterization of implementability.

Lemma **3.** *An approval rule $y : \mathcal{T} \to [0, 1]$ is implementable if and only if it is nondecreasing.*

If an approval rule is implementable, then there is a standard of proof that implements it so that the menu is IR. Let $\mathcal{I}$ denote the set of implementable approval rules $y : \mathcal{T} \to [0, 1]$.

3.3 Regulator's Information and Objectives

The regulator is evaluated on how large an error she makes in her approval decisions, where errors scale with her extent of approval. The regulator faces ambiguity over Θ and attempts to minimize worst case payoff over this space. As seen in Section 2, we can represent such a regulator as minimizing *worst case regret*. How the regulator designs approval rules depends on the criteria on which she is evaluated. In general, the regulator is evaluated both on the basis of errors in decision making and the costs of regulation imposed on firms. She compares the realized outcome to the outcome had she been informed of the firm's type *a priori* and not asked the firm to generate costly evidence.

Formally, the regulator's worst case evaluation for an approval rule $y \in \mathcal{I}$ is:

$$R(y) \equiv \sup_{\theta \in \Theta} \max\{v(\theta), v_0\} - [y(\tau(\theta))v(\theta) + (1 - y(\tau(\theta)))v_0] + h(\theta)n(\tau(\theta)) \tag{1}$$

where $h : \Theta \to \mathbb{R}_+$. h captures now negatively the regulator is evaluated ex-post for making a type θ firm expend effort $n(\tau(\theta))$ to prove its efficacy. Depending on how the regulator is being evaluated, $v(\cdot)$ may encode a concern with only the quality of the firm's innovation relative to the status quo; or it might also account for how many consumers were actually negatively affected by the regulator's decision relative to the best action in hindsight.

The regulator chooses a $y \in \mathcal{I}$ and corresponding standard of proof n which implements y with $n(\underline{\tau}) = 0$ to solve:

$$\inf_{y \in \mathcal{I}} R(y) \tag{2}$$

Section 4 analyzes the special case where $h = 0$: the regulator is judged on *consumer protection*. Section 5 considers the case where $h > 0$: the regulator is judged on *consumer protection* and *firm approval costs*.

3.4 Discussion of Assumptions

Firm's Private Information The assumption that the firm knows its benefit of approval exactly can be relaxed. b can instead be interpreted as the firm's *expected* benefit, and Appendix A shows that a risk-neutral regulator still behaves as if she minimizes worst-case *expected* regret when outcomes are stochastic.

The assumption that the firm knows q is more appropriate if this is easily learned in the course of developing the product or conducting tests. For example, an AV manufacturer knows how frequently its product requires human intervention and would therefore learn little from logging miles without intervention.[13] A company producing drugs for treating heart-disease may only be required by the FDA to show short term test results exhibiting how well the drug reduces bad cholesterol. The company may be much more knowledgeable about how it fares in terms of this surrogate endpoint than it would be about the genuine efficacy of its drug.

Regulator's Actions The firm is assumed to prefer more approval to less, regardless of its type. The interpretation is that by changing y, the regulator changes the number of potential customers that the firm has access to in a way such that those with and without access are representative of the full population.[14] For example, the regulator may allow immediate passage of self driving cars in some states.

The regulator may also achieve a similar effect by delaying approval. If a new batch of potential customers arrive each period (e.g., patients falling ill to a certain condition) and each has the same distribution of preferences, delay simply reduces the firm's market size without changing the firm's pricing and quantity decisions. Alternatively, the population of potential consumers may stay the same, and delay causes the firm's payoffs to be discounted at some rate (alternatively, some change in the environment arrives at some rate and renders the proposed class of innovations obsolete).[15]

[13] This is policy, for example, in California: "California requires AV testing companies operating in California to document and report miles driven as well as the number of times when a human driver, when present, had to retake control of the vehicle" (Brown et al., 2018).

[14] A regulator may have power to limit supply more directly by placing quantity caps on the innovating firm, rather than restricting market exposure. In contrast to market exposure, the firm need not strictly prefer higher quantity caps to lower ones once the constraint exceeds the monopoly optimal quantity chosen by the firm. This case again can be dealt with a slight adjustment to the existing model.

[15] The model currently assumes that the regulator's benefit v and the firm's payoff b scale the same way with y, but easily extends to the case where they scale differently.

4 Consumer Protection

This section considers a case of the model where the regulator is evaluated solely based on consumer interests. Section 4.1 gives a simple example to illustrate the model and how optimal rules are constructed. Section 4.2 solves the problem for a regulator evaluated only on the basis of consumer protection. Section 4.3 describes how certain naturally arising implementability considerations in the consumer protection problem shape regulator-optimal approval rules.

4.1 Partial Disclosure Example

An agency regulating vehicle safety, like the NHTSA, concludes that their value for approving a particular automated vehicle model lies in $[\underline{v}, \overline{v}]$. The firm's type θ is a vector of characteristics of the car and the firm (physical design features of the car, cost of manufacturing and employing workers, marketing strategy etc.), which completely determines the regulator's value for approval, $v(\theta)$. Suppose that $v(\theta) = \alpha q(\theta)$ for some $\alpha > 0$, where $q(\theta) \in [\underline{q}, \overline{q}]$ is the expected fraction of trips where the AV's algorithm produces no errors. Assume that the benefit of approval to the firm is independent of its type and normalized to 1.[16]

The car manufacturer collects data to prove the safety of its cars in a variety of natural situations at a cost $k(\theta) \in [\underline{k}, \overline{k}]$ per trip. Let $c(\theta) \equiv \frac{k(\theta)}{q(\theta)}$, the expected cost of running enough trips to produce one incident-free trip.

The firm cannot fabricate evidence of successful trips its drivers have taken, but it can 1) selectively disclose data so as not to reveal those involving errors or requiring human intervention and 2) intentionally create trips with errors at no cost.

The regulator evaluates her choice of allocation rule x by the size of the type I or type II mistakes she makes, in the worst case. That is, given a benchmark of v_0 and $\hat{\rho}_0$ reported failed trials and $\hat{\rho}_1$ reported successful trials by a firm with a product of quality v, the regulator has an ex-post payoff given by

$$R^C(v_0, v | x, \hat{\rho}_0, \hat{\rho}_1) = \max\{(1 - x(\hat{\rho}_0, \hat{\rho}_1))(v - v_0), x(\hat{\rho}_0, \hat{\rho}_1)(v_0 - v)\}.$$

For a given allocation rule x, let (y, n) denote the induced approval rule and standard of proof. A firm of type θ optimally chooses to produce $n(\frac{1}{c(\theta)})$ successes for approval level $y(\frac{1}{c(\theta)})$, giving it a payoff of $y(\frac{1}{c(\theta)}) - n(\frac{1}{c(\theta)}) \cdot c(\theta)$.

The regulator effectively chooses an approval rule y to solve

$$\min_{y \in \mathcal{I}} \max_{q,k,v} \max\{(1 - y(\frac{q}{k}))(v - v_0), y(\frac{q}{k})(v_0 - v)\}.$$

Suppose the interpretation of y is the fraction of districts within a country for which AVs are approved, with $y = 0$ meaning that the innovation is rejected (at least until the

[16] One may think, for example, that while the variation in safety specifications and accident rates matter to the regulator, the differences in safety are imperceptible to consumers within the range $[\underline{v}, \overline{v}]$.

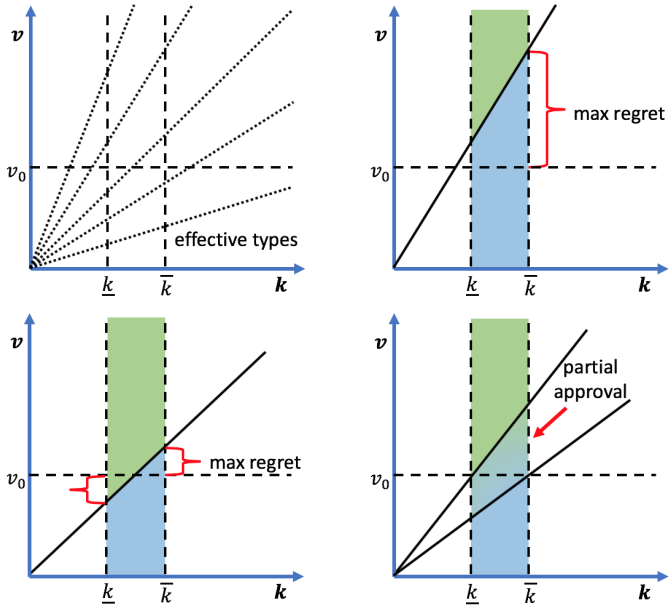

Figure 2: The type space is the column between $\underline{k}$ and $\overline{k}$. The upper left figure shows iso effective type curves, which are the intersection of the type space and the rays emanating from the origin. The figures on the top right and bottom left depict participating (green) and nonparticipating (blue) types under different cutoff approval rules. The fully optimal screening contract is depicted on the bottom right. It consists of a region of partial allocation where effective types smoothly increase from full rejection to full approval.

revelation of the actual quality of the AV model) and $y = 1$ meaning that the innovation is approved nationwide. Consider first the regulator's choice of approval rule within the space of cutoff rules: approve the product nationally if at least $n^* \in \mathbb{R}_+$ successes are shown and rejects otherwise. For any standard of proof n^*, the firm will engage in the task of evidence provision only if expected payoff of doing so is positive, i.e., if $1 - n^* \cdot c(\theta) \geq 0$. This means the firm provides the evidence necessary for immediate approval whenever $q(\theta) \geq n^* \cdot k(\theta)$. Figure 2 shows the projection of the firm's type space Θ onto (regulator outcome)$\times$(marginal cost) space. A ray of slope $\alpha \cdot n^*$ divides types who would obtain approval and those for whom providing the requisite evidence is prohibitively costly.

The largest mistakes from any such cutoff rule come from one of two types of the firm. The first type has very high costs running trips and opts out of the regulator's test. To the regulator's dismay, this happens even though the firm has a product safe enough to push the regulator's value of approval over v_0. The other type of firm produces an automated vehicle much worse than the status quo but enjoys a low cost of generating trips. It games the regulator's approval standard by suppressing failed test runs from surfacing. The optimal choice of n^* is therefore one which equalizes the largest size of either type of error.

But cutoff rules are typically sub-optimal. Optimal rules in general take the form of a menu where sufficiently low types never get approved, and sufficiently high types are

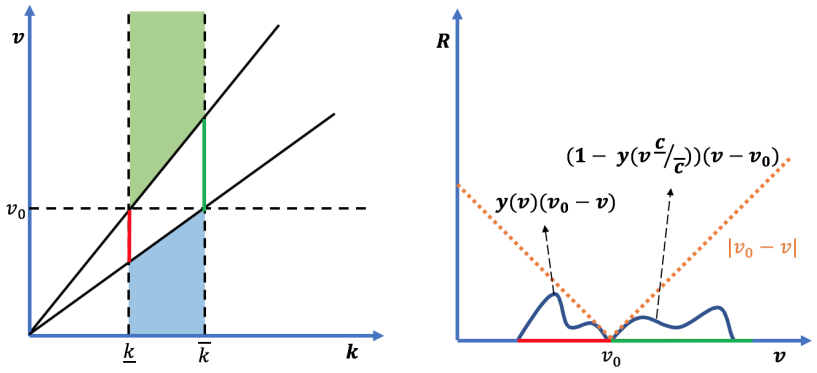

Figure 3: On the left, the white area between the green and blue regions shows types for whom the regulator cannot avoid making approval errors. The firm types highlighted in red are those who have lowest marginal cost of running trials among those types she prefers to fail. The types highlighted in green are those who have the highest costs among types she prefers to pass. On the right, the regulator's outcome v of passing these red and green firm types are represented in the x axis, and the regret of making the wrong approval decision with these types is plotted in orange. The solid blue line plots the regret of an arbitrary approval rule y for these same types.

approved immediately, and intermediate types receive partial approval.

In this example, a particular regret-minimizing approval rule y^* can be constructed by applying Lemma 1 to each iso-effective type subset of Θ. Higher effective types get more approval in exchange for showing more evidence. Such an approval rule is pictured on the bottom right of Figure 2. Even if the regulator could not commit to an approval rule, if it could observe the firm's effective type (while the firm's type remains private information), it would still choose to give $y^*(\tau)$ approval to a firm of type τ.

This optimal approval rule is not unique, and there is a solution other than y^* that is of interest. Since an approval rule is evaluated at its worst case effective types, the level of approval for other effective types may be altered without raising the max regret of the original approval rule. Indeed, in this example, Figure 3 and Figure 4 illustrate how to graphically construct a solution where 1) all sufficiently low effective types achieve the max regret, and 2) all higher types have full approval. Under this "constant regret, full approval" approval rule, the worst case error that a regulator can commit on any effective type of the firm is always a type I error. While this example shows that there can a multiplicity of solutions to the consumer protection problem, the solution in Figure 4 will emerge in the limit as the regulator's concern about producer costs approaches 0, as Section 5 will make clear.

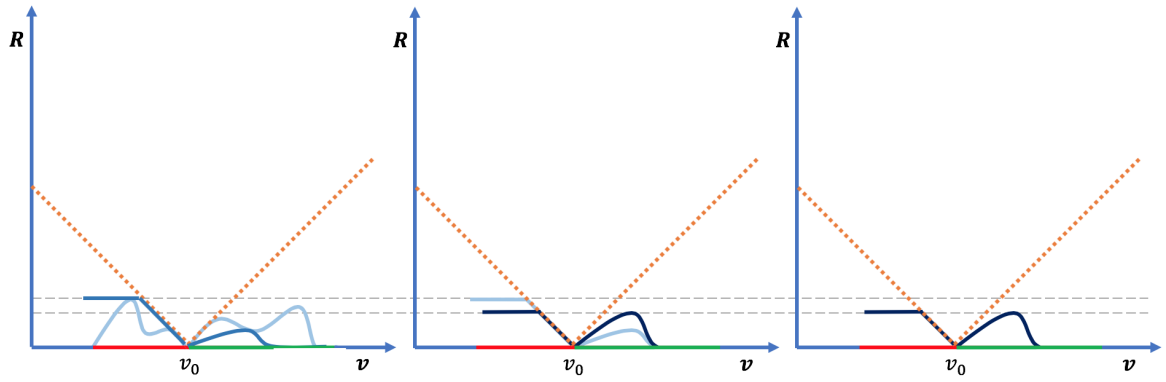

Figure 4: The regret on a "red firm" is increased by giving it more approval. But this reduces the regret for the corresponding "green firm" with the same effective type. Taking the regret profile of an arbitrary approval rule, we can increase the regret on red types so as not to increase the max regret of the contract (left). This regret profile corresponds to an approval rule which has a constant regret on red types and almost everywhere lower regret for the green firms. From this regret profile, we now have leeway to move to another approval rule that has a lower constant regret level of regret on red firms but higher regret on green firms (middle). If the highest regret on green firms is the same as the highest regret on red firms, the corresponding approval rule must be optimal (right). We could have started this procedure by moving to a constant regret level on the green firms, and that would have yielded a different optimal approval rule.

4.2 Regret Minimizing Approval Rules

A solution to the regulator's problem (2) is called a *regret minimizing approval rule*. Let $R^C(\cdot)$ denote the objective function mapping approval rules to worst case regret in the case of pure consumer protection, i.e., $h = 0$. I henceforth assume that Θ is rich enough that the regulator cannot avoid every error:

ASSUMPTION **2.** *There exist* $\theta, \theta' \in \Theta$ *such that* $\tau(\theta) = \tau(\theta')$ *and* $v(\theta) > v_0 > v(\theta')$. *There exist* $\theta, \theta' \in \Theta$ *such that* $\tau(\theta) > \tau(\theta')$.

In particular, Assumption 2 implies that all cutoff rules are suboptimal.

Let $\Theta_\tau \subset \Theta$ denote the set of firm types with effective type τ. Let $v(\Theta_\tau) \subset \mathbb{R}_+$ denote the image of Θ_τ under v. Finally, let $l(\tau) \equiv \inf v(\Theta_\tau)$ and $m(\tau) \equiv \sup v(\Theta_\tau)$. Define the regret arising from an allocation rule conditional on a given effective type as follows:

$$R^C_\tau(y) \equiv \sup_{\theta \in \Theta_\tau} \max\{v_0, v(\theta)\} - (1 - y(\tau))v_0 - y(\tau)v(\theta).$$

I restrict attention to those approval rules which are *undominated*: y is said to dominate $\hat{y}$ if the former gives the regulator a weakly better payoff in all states of the world (i.e., for all realizations of the firm's type, $\theta \in \Theta$, $y(\theta)v(\theta) + (1 - y(\theta))v_0 \geq \hat{y}(\theta)v(\theta) + (1 - \hat{y}(\theta))v_0$) and a strictly better payoff in some state. $\hat{y}$ is undominated if no y dominates $\hat{y}$.

An approval rule y is *admissible* if,

1. $y(\tau) = 1$ if $l(\tau) \geq v_0$ and $m(\tau) > v_0$
2. $y(\tau) = 0$ if $m(\tau) \leq v_0$ and $l(\tau) < v_0$

The first criterion, for example, says an effective type is passed when there is a possibility of an upside but no possibility of a downside. The following lemma characterizes undominated approval rules in this setting.

LEMMA **4.** *An approval rule $y : \mathcal{T} \to [0,1]$ is undominated if and only if it is admissible.*

Let the approval rule y^* be defined such that $y^*(\tau) = \arg\min_{p\in[0,1]} R^C_\tau(p)$ for all τ; Lemma 1 gives the expression for $y^*(\tau)$.

PROPOSITION **1.** *If y^* is increasing, it minimizes regret among all undominated and implementable approval rules.*

The proof in Appendix E shows a stronger statement, characterizing all undominated, regret-minimizing approval rules. Note that one sufficient condition for y^* to be increasing is that $l(\cdot)$ and $m(\cdot)$ are increasing.

Suppose a set of firms have types drawn from some full support distribution on Θ, and that these firms participate in the regulator's approval process. By Lemma 3 and Proposition 1, approval levels depend only on and are increasing in effective type. On the other hand, a firm A with lower costs or higher benefits of approval than firm B can have more delay if it is of a lower effective type. This accords with the findings of Olson (1997) that upon controlling for firm and drug characteristics, the correlation between FDA approval delays and the ratio of research costs to ex-post revenue was significant, while the correlation to costs alone was insignificant. This fact indeed suggests that regulators see information beyond reported findings in firm costs and expected revenues, consistent with a model of experimental manipulation.

4.3 Interpreting Intermediate Evidence

Monotonicity of $l(\cdot)$ and $m(\cdot)$ is violated in many consumer protection contexts.

Imperfect Markets or Information A firm may have a low benefit of approval, and therefore a low effective type, for one of two reasons. Its product may be of obviously low quality and fail to reach a wide population. Alternatively, the firm's benefit may be low simply because it makes a niche product, which is still effective and safe. Given the market's ability to partially self-regulate, raising standards of proof high enough to weed out low quality products may not be worth the cost of precluding such niche products from passing. The issue however may arise with products of seemingly intermediate quality. If indeed they turn out to be worse than existing substitutes, they may nevertheless reach a wide audience upon approval.

The same issues can arise when consumers are perfectly informed if the regulator's interests and consumer interests are divergent (e.g., the former cares about quality, the latter care about quality and price). See Appendix D for a numerical example.

Behavioral Responses It may be the case that AVs that approach the regulator's desired standard of safety, measured by rate of incident-free trips, may be more dangerous than those with much higher incident rates. A driver may be much more vigilant in a vehicle known to require frequent human intervention than in one that rarely fails. But this compensating behavioral response may make the latter vehicle type more dangerous in terms of fatal crashes. In the best case, the regulator may still think the lower incident rate car is the safer one. But in the worst case, approval may lead to more accidents than possible with a higher incident rate car.

These examples suggest scenarios where l is initially decreasing and subsequently increasing, while m is nondecreasing. Denoting τ' as the effective type at which the slope of l changes sign. Consider the approval rule,

$$y_0 = \begin{cases} y^*(\tau'), \text{ if } \tau \in [\underline{\tau}, \tau'] \\ y^*(\tau), \text{ if } \tau > \tau' \end{cases}$$

which pools effective types in $[\underline{\tau}, \tau']$. This is a regret minimizing approval rule, and $R^C_{\tau'}(y^*(\tau'))$ is the max regret. To see why, note that the interval $[l(\tau'), m(\tau')]$ contains $[l(\tau), m(\tau)]$ for all $\tau < \tau'$. This means any action that the adversary can take to best respond to the regulator's choice of delay at $\tau < \tau'$, it can also take at τ' and produce weakly larger errors.

Approval rule y_0 under-serves low effective types relative to the (unimplementable) rule, y^*. A corollary in Appendix B implies that when $R(y_0) = R_{\tau'}(y_0)$, all optimal rules do the same. Upon reviewing the submitted evidence, a regulator awards firm effective types in $[\underline{\tau}, \tau']$ *less approval* than she would like. Committing to doing so helps her avoid making even costlier mistakes on the firms that show intermediate levels of approval.

5 Consumer Protection and Regulatory Costs

In addition to the ex-post correctness of their actions, regulators are also frequently evaluated on the basis of the costs they create for regulated industries, once they are publicly revealed. A common criticism of the FDA is that the costs they create for firms to approve their products are in part reflected in drug prices paid by consumers.

Section 5.1 shows that unlike the case of pure consumer protection, regulators concerned with costs of evidence provision commit to give *more approval* to firms showing low levels of evidence than what they would like. They even give some level of approval away without seeing any evidence. This can be interpreted as setting up product testing

regulatory sandboxes, which is a common practice in fintech regulation. Section 5.2 characterizes regulator optimal rules when additional assumptions are made about the type space. Delegation sets might preclude the option of full approval, and optimal approval rules may be unique unlike the case where evaluations only measure consumer protection. Section 5.3 discusses the interpretation of these results.

5.1 Internalizing Evidence Production Costs

A regret minimizing regulator compares the realized costs of evidence provision and the correctness of the approval decisions to the outcome realized she would have realized had she known the firm's type prior to engaging the latter party in the task of costly information provision.

The regulator chooses an approval rule $y \in \mathcal{I}$[17] and corresponding standard of proof n which implements y with $n(\underline{\tau}) = 0$ to minimize:

$$\begin{aligned} R(y) &= \sup_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} \max\{v(\theta), v_0\} - [y(\tau(\theta))v(\theta) + (1 - y(\tau(\theta)))v_0] + h(\theta)n(\tau(\theta)) \\ &\equiv \sup_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} R_\tau(y, \theta) \\ &\equiv \sup_{\tau \in \mathcal{T}} R_\tau(y), \end{aligned}$$

where $h : \Theta \to \mathbb{R}_+$ is a strictly positive valued function representing the costs of producing successful evidence. h may differ from c: the regulator may be be evaluated, for example, on the basis of the environmental and traffic costs of running extensive AV trials, and not the firm's incurred costs of evidence production directly.

PROPOSITION **2.** *Suppose that $l(\tau)$ is nondecreasing. An optimal approval rule for a regret minimizing regulator who internalizes the firm's welfare 1) awards strictly positive level of approval is given even when no evidence is shown, 2) attains max regret at the lowest effective type.*

Proof. Take any implementable y such that $R(y) > R_{\underline{\tau}}(y)$. Note that $y(\underline{\tau}) \neq 1$. Otherwise, implementability would require that $y(\tau) = 1$ and $n(\tau(\theta)) = 0$ for all τ; this coupled with the fact that $l(\cdot)$ is nondecreasing would imply that $R(y) = R_{\underline{\tau}}(y)$, a contradiction.

Choose $\epsilon > 0$ small enough so that $R_{\underline{\tau}}(y') < R(y)$, where

$$y' \equiv \begin{cases} y(\underline{\tau}) + \epsilon, & \text{if } y < y(\underline{\tau}) + \epsilon \\ y, & \text{if } y \geq y(\underline{\tau}) + \epsilon \end{cases}$$

[17]Note the earlier notion of admissibility is dropped. In fact, all admissible rules will typically be suboptimal when the regulator internalizes the firm's cost of evidence.

and $n'(\tau)$ implements y' with $n'(\underline{\tau}) = 0$. Consider τ such that $y(\tau) < y'(\tau)$. Note that $n'(\tau) = 0$, so

$$\begin{aligned} R_\tau(y') &= \max_{\theta \in \Theta_\tau} \max\{(1 - y'(\tau))(v(\theta) - v_0), y'(\tau)(v_0 - v(\theta))\} \\ &= \max\{(1 - y'(\tau))(\max_{\theta \in \Theta_\tau} v(\theta) - v_0), y'(\tau)(v_0 - l(\tau))\}. \end{aligned}$$

If $R_\tau(y') = y'(\tau)(v_0 - l(\tau))$, then $R_\tau(y') \leq R_{\underline{\tau}}(y') < R(y)$, since $l(\cdot)$ is nondecreasing. And if $R_\tau(y') = (1 - y'(\tau))(\max_{\theta \in \Theta_\tau} v(\theta) - v_0)$, then $R_\tau(y') < (1 - y(\tau))(\max_{\theta \in \Theta_\tau} v(\theta) - v_0) \leq R_\tau(y) \leq R(y)$.

Now suppose τ is such that that $y(\tau) = y'(\tau)$. Then $n'(\tau) = \tau y'(\tau) - \underline{\tau}(y(\underline{\tau}) + \epsilon) - \int_{\underline{\tau}}^{\tau} y'(t)dt < \tau y(\tau) - \underline{\tau} y(\underline{\tau}) - \int_{\underline{\tau}}^{\tau} y(t)dt = n(\tau)$. Since $h > 0$, $R_\tau(y, \theta) > R_\tau(y', \theta)$ for every θ, so $R_\tau(y') < R_\tau(y) \leq R(y)$.

Therefore $R(y') < R(y)$, so the chosen y is not optimal. This shows that if y is optimal, $y(\underline{\tau}) > 0$ and $R_{\underline{\tau}}(y) = R(y)$. □

Firms may decide to show no evidence. But rather than reject such firms outright, the regulator gives them limited approval, like in a regulatory sandbox. In particular, this is also true when $m(\underline{\tau}) < v_0$. Partially approving even those firms who are guaranteed to generate outcomes worse than the status quo reduces the burden of evidence that higher effective types would need to show in order to secure more approval.

5.2 Optimal Approval Rules: Interval Type Spaces

The solution to the regulator's problem can be fully characterized with the help of some assumptions on the range of costs and benefits to the regulator.

Assumption **3** (Interval Type Space)**.** *For every τ, there exist $\theta', \theta'' \in \Theta_\tau$ such that $c(\theta') = \underline{c}$ and $c(\theta'') = \overline{c}$. The regulator's benefit of approval is within an interval of the firm's benefit of approval: $v(\theta) \in [b(\theta) - d(\tau(\theta)), b(\theta) + D(\tau(\theta))]$ for all $\theta \in \Theta_\tau$, for some differentiable functions $d, D \geq 0$, where d is nonincreasing and D in nondecreasing. Moreover, $v(\theta') = b(\theta') - d(\tau(\theta'))$ and $v(\theta'') = b(\theta'') + D(\tau(\theta''))$.*

These assumptions, for instance, nest the case where all types of the firm are equally likely to pass tests but have different benefits in obtaining approval. In the case of AVs, when the probability of human intervention is very low for any type of firm, the biggest source of variation in the difficulty of logging successful trips arises from differences in operational costs. Moreover, if accident statistics are public and newsworthy (as they tend to be for new technologies), differences in firm benefits to approval are likely to drive differences in willingness to show evidence. This case is also covered by Assumption 3. The regulator minimizes the following objective with respect to $y \in \mathcal{I}$:

$$\sup_\tau \max_{\theta \in \Theta_\tau} \max\{v_0, v(\theta)\} - [y(\tau)v(\theta) + (1 - y(\tau))v_0)] + c(\theta)n(\tau)$$

Consider the special class of approval rules which have constant regret on one region, followed by constant allocation on the next.

DEFINITION **1.** *y is a **constant regret, constant allocation (CRCA)** approval rule if there is some $\tau' \in [\underline{\tau}, \overline{\tau}]$ that partitions the effective type space into two regions: $R_\tau(y) = R(y)$ on $[\underline{\tau}, \tau']$, and $y(\tau) = y(\tau')$ on $(\tau', \overline{\tau}]$. Let Y_{CARA} denote the set of approval rules that are constant regret, constant allocation for some level of regret, R.*

PROPOSITION **3.** *When the regulator is judged on consumer protection and firm approval costs, a CRCA optimal approval rule exists. If $D(\overline{\tau}) < v_0$, this rule is the unique solution to the regulator's problem.*

The main step is to show that for any implementable allocation rule, y, there exists an approval rule $y' \in Y_{CARA}$ such that, $R(y') \leq R(y)$. For ease of exposition, I will say hasten and delay approval to refer to increasing and decreasing the extent of approval.

The monotonicity conditions on d and D ensure that the if the regulator could directly observe effective types, she approve higher types sooner. To minimize firm costs, she would want to hasten approval for low effective types so long as this does not increase max regret of the overall contract. This reduces the burden of evidence for higher types who would only need to show a little more evidence to separate themselves from lower types and be approved even sooner.

Given any approval rule y with $R(y) = R$, the regulator can hasten the process for low effective types up to the point where the regret incurred for each effective type is either R or the extent of approval is $y(\overline{\tau})$. This construction begins with the smallest effective type and builds upwards, since hastening approval for lower effective types reduces the regret for any level of approval for higher effective types.

Letting $\tilde{y}$ be the solution to the integral equation $R_\tau(\tilde{y}) = R$, the construction defined earlier is the CRCA approval rule $y' = \min\{\tilde{y}, y(\overline{\tau})\}$, which is non-decreasing under the assumed conditions.

Consider the regret on the highest effective type. Since $R_{\overline{\tau}}(y) \leq R$, it certainly the case that $R_{\overline{\tau}}(y') \leq R$. This is because $y(\overline{\tau}) = y'(\overline{\tau})$, but $y' \geq y$ on $[\underline{\tau}, \overline{\tau})$. So under approval rule y', the highest effective type gets the same level of approval as in y, but potentially has to show less evidence for it.

Next note that $R_\tau(y') \leq R$ on the portion of the contract where y' is constant, as pictured in Figure 5. In this region, regret initially dips in the region where the worst case error is of type I.[18] The size of type I errors is decreasing since $l(\tau)$ is increasing. Thereafter, the worst case error becomes type II, and this increases as $m(\tau)$ increases. This means regret drops below R on this region but never rises above $R_{\overline{\tau}}(y')$, which is also less than R. The search for optimal rules can be therefore be restricted to the space of CRCA rules.

[18] A bulk of the formal proof in Appendix E shows that the worst case error on this region cannot initially be of type II.

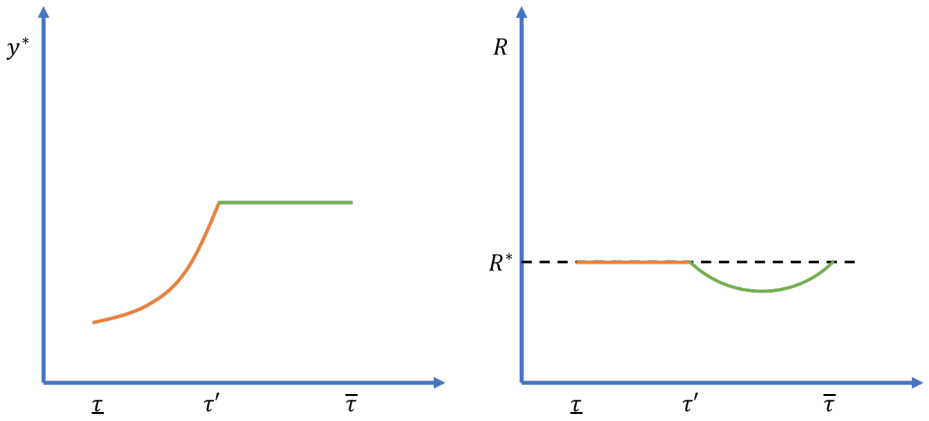

Figure 5: On the left is an example optimal approval rule y^*, which has constant regret on the orange portion and is constant on the green portion, as shown on the right. The optimal rule has the feature that max regret is achieved at $\overline{\tau}$.

In the space of CRCA approval rules, it is easy to check for further improvements.

If $R_{\overline{\tau}}(y') < R$, then on the region where y' was increasing, let y'' be the solution to an integral equation $R_{\tau}(y'') = R - \epsilon$, where $\epsilon > 0$ is small, and let y'' be constant afterwards. $R_{\overline{\tau}}(y'')$ may now be higher than $R_{\overline{\tau}}(y')$ while still being below $R - \epsilon$. This means $R(y'') < R(y')$.

Similarly, we can shift the point where the CRCA switches from being constant regret to constant allocation to see if this reduces the regret on $\overline{\tau}$. If it does, this would suggest again moving to a CRCA contract where the constant regret portion is of lower regret. In an optimal approval rule, there is typically no slack between the regret on the highest effective type and the regret over the entire approval rule , as pictured in Figure 5.

A CRCA approval rule pools high effective types. In particular, note that the condition pinning down τ', the effective type after which all types are pooled at the same delay level, does not depend on $l(\overline{\tau})$, which may be greater than v_0. In this case, the best decision if the regulator could observe effective types would be to approve type $\overline{\tau}$ firms immediately. But the optimal menu may exclude the option of immediate approval for any level of evidence: there is inefficient distortion at the top. In such cases, a regulator who can partially approve innovation for some markets waits for the state of the world to be exogenously revealed before upgrading the firm to full approval or banning the product.

5.3 Permissionless Innovation and Strict Precaution

It is often argued that when firms can provide selective evidence, regulators can do little to determine whether or not an innovation should be approved:

> The task of establishing facts about new technology may be made difficult by the lack of an adequate sample or other reliable data on the effects of new technol-

> ogy[...]other facts may be—to quote Donald Rumsfeld— "unknown unknowns." We simply lack the experience or imagination to predict what negative possibilities may be associated with a piece of new technology. In this respect, the "relevant facts" that form the basis of regulation are never going to be obvious or settled. The regulation of any disruptive new technology is always going to be reactive and based on an uncertain and politicized factual basis. (Fenwick et al., 2016, p.574)

These arguments are frequently followed up with prescriptions (depending on the interests of the party) that regulators follow one of two modes. Innovating technology firms (e.g., Uber) often ask regulators to follow a model of "permissionless innovation", meaning that they should only regulate innovations ex-post if they turn out to be harmful (Thierer, 2016). Other parties (e.g., consumer groups) may instead want regulators to ban products until evidence guarantees their safety or efficacy. This mode of regulation follows a strict interpretation of what has come to be known as the "precautionary principle". European regulators sometimes appeal to this principle to justify decisions, for example, like banning sale of food products and crop with GMOs in many member EU countries until incontrovertible exogenous evidence arises proving their safety.

When ambiguity averse regulators are evaluated on the correctness of their decisions ex-post and the firm's incurred regulatory costs, they may voluntarily choose approval rules that share features of both permissionless innovation and strict precaution. Regulators may award partial approval in the form of regulatory sandboxes even without having seen prior evidence; and they may wait for the arrival of incontrovertible evidence before providing full approval, even if firms have the ability to credibly signal their superiority to the status quo at the time of evidence provision.

6 Discussion

Section 6.1 clarifies the role of partial approval in the regret minimizing approval rules by considering regulators who either do not face ambiguity or are unconcerned with type II errors. Section 6.2 discusses why the setting of screening through costly manipulable evidence might be a reasonable model of many regulatory agencies.

6.1 Role of Partial Approval

Partial approval, in the form of delays or regulatory sandboxes, can occur for many reasons outside of the scope of this model. For example, partial approval may hasten the arrival of information about innovation quality.

However, as discussed in Section 2.3, regulator incentives can account for partial approval or delay when alternative explanations are inadequate, or may reinforce delays stemming from information considerations in other cases.

Section 2 showed that when the only form of communication between firms and the regulator was cheap talk, both ambiguity over product characteristics and concern for type II errors (which Carpenter (2004) argues have become more visible and a bigger concern for regulators than in the past[19]) were necessary features for explaining partial approval. Here, I show that the same is true in the model with costly evidence generation and selective disclosure.

To this end, define a *cutoff policy* as an approval rule y for which there exists a τ^* such that $y(\tau) = 1$ for $\tau \in (\tau^*, \overline{\tau})$ and $y(\tau) = 0$ for $\tau \in (\underline{\tau}, \tau^*)$.

6.1.1 Innovation Approval Without Ambiguity

Consider the case where the regulator a prior F over Θ and is again risk neutral.

Suppose that the marginal distribution of F over $\mathcal{T}$ admits a continuous density function, f_τ. Let $E[g(\theta)|\tau]$ denote $\int_{\theta \in \Theta_\tau} g(\theta) dF(\theta|\tau)$ for any function $g : \Theta \to \mathbb{R}$. The objective of the Bayesian regulator is given by:

$$V_B(y, n) = \int_{\underline{\tau}}^{\overline{\tau}} (E[v(\theta)|\tau]y(\tau) + v_0(1 - y(\tau)) - E[h(\theta)|\tau]n(\tau))dF_\tau(\tau),$$

and she chooses $y \in \mathcal{Y}$ and an n_y that implements it with $n_y(\underline{\tau}) = 0$ to solve:

$$\max_{y \in \mathcal{I}} V_B(y, n_y).$$

The next proposition shows that a Bayesian regulator never strictly benefits from using partial approval.

PROPOSITION **4.** *The Bayesian regulator's optimal approval rule can be implemented as a cutoff policy.*

The intuition for the result is similar to why an optimal auction to a single buyer can be implemented as a fixed price mechanism.

6.1.2 Innovation Approval Without Concern For Type II Errors

Suppose now that the regulator faces ambiguity over types θ but is never punished for type II errors. That is, she behaves as if she maximizes minimum absolute payoffs instead:

[19] "Before the 1980s it was rare for anyone outside of clinical or academic circles to criticize the FDA for delay. Put differently, few in the media or in Congress were complaining of the agency's Type II errors.

Today the situation is much different, and the political power of patients matters as much as or more than the political power of firms. The best evidence for this proposition comes in two strategies that are now widely adopted by pharmaceutical firms: (1) Firms themselves have in the past six to eight years created, fostered, and subsidized a number of patient advocacy groups; and (2) firms regularly seek alliances with patient advocates in pressing the case for priority status, accelerated approval, or simply approval before the FDA." (Carpenter, 2004).

$$\max_{y\in\mathcal{I}} V_M(y) \equiv \max_{y\in\mathcal{I}} \min_{\theta\in\Theta} y(\tau(\theta))v(\theta) + (1 - y(\tau(\theta)))v_0 - h(\theta)n(\theta)$$
$$\equiv \max_{y\in\mathcal{I}} \min_{\theta\in\Theta} J(y,\theta)$$

PROPOSITION **5.** *An optimal, undominated approval rule of a max-min regulator who faces ambiguity over Θ and internalizes the firm's welfare is the following cutoff policy:*

$$y^*(\tau) = \begin{cases} 0, \text{ for } \tau < \tau^* \\ 1, \text{ for } \tau \geq \tau^* \end{cases}$$

where $\tau^* = \inf\{t \in \mathcal{T} | l(t') \geq v_0, \forall t' \geq t\}$. If the regulator is judged on consumer protection and firm costs, this optimal approval rule is unique.

The difficulty of the approval threshold is chosen so that even the worst type of firm that finds it worthwhile to submit evidence is at least as good as the status quo from the regulator's standpoint.

6.2 Costly Standards of Proof as a Second-Best Solution

The main model assumes that (1) firms communicate their private information by choosing how much evidence to generate and what to show[20], and (2) that the main tool of the regulator is the ability to choose the extent of approval to award. Section 2.3 discussed direct suggestive evidence that this is a reasonable model of the FDA review process.

But ideally, a regulator would be able to circumvent the issues of manipulation altogether by learning about an innovation through direct testing or by inducing truthful reporting by punishing misreporting. Here I consider some reasons why these first best solutions may not be feasible in practice.

I use automated vehicles and new drug approval as key examples in discussing these issues, though they arise with other innovations as well (e.g., medical devices or new fin-tech products).

Direct Testing Regulators need no solicit information if it is easy for them to run tests directly. In the case of drug approval, it is well known that tests are costly and sophisticated.

It also turns out that testing the safety standards of an automated vehicle is significantly more challenging than testing non-automated vehicles.[21] Currently, new car

[20] Appendix C considers a model of communication in which evidence is directly informative, which may be more realistic in cases where the firm is limited in its ability to hide negative evidence.

[21] One alternative is to have manufacturers pay third-party auditors to conduct these tests on their behalf and present results to the relevant regulatory agency. However, delegating data collection to a third party may simply increase the firm's cost of partial disclosure rather than do away with the

models are tested with a battery of standardized collision scenarios.[22] Gaming these performance metrics without having a truly safe vehicle is difficult, and conducting these tests is straightforward for third-parties. Testing the safety of an automated vehicle, on the other hand, involves checking how the automated decision-making performs in a plethora of natural driving conditions. Subtle differences in road and traffic conditions which are imperceptible to human drivers may trigger very different responses by an automated vehicle. Mistakes made by other human drivers are rare and varied as well. This makes the data collection task for verifying safety extremely onerous. It takes 11 billion miles driven, by one estimate, to say with 80% confidence that an automated vehicle model is 90% safer than human driven cars, and feasible tests to bring this number down is an active area of research (Zhao, 2016). Currently, it appears infeasible for regulatory agencies to test vehicle quality directly.

Regulating Product Features When considering new innovations, regulators face difficulties in applying command-and-control type product design regulations, which constitute many laws aimed at promoting consumer safety and welfare. The U.S. Environmental Protection Agency, for example, mandates through the Clean Air Act and Clean Water Act that firms adopt certain technologies to limit emissions, such as scrubbers on smoke tanks. Or the NHTSA requires cars have seat belts and other safety features to be street legal.

A regulatory body would likely not have the specialized knowledge required to make appropriate suggestions regarding what fail-safes to include in the algorithm of an automated vehicle, what hardware specifications should be met, etc. (And of course, suggesting changes to chemical compositions of drugs is wholly inappropriate.)

Formalizing Experimentation Process Another suggestion is to put in place various formal requirements, like pre-registration of experiments, that limit the ability of a firm to manipulate evidence.

It is first important to note that the process of drug approval is now fairly formalized, and hiding trials is significantly more difficult for a pharmaceutical company today than ever before. Still, other types of manipulation are still possible and difficult to catch. Examples include ending trial too early when positive results are recorded, using favorable surrogate end-points, and selecting samples that inflate experimental success rates. Appendix C shows that the forces studied in this model are still applicable in a setting with limited manipulation.

agency problem entirely. Moreover, it is unlikely that auditing services with the narrow expertise in testing automated vehicles would be around at the time manufacturers seek approval. For one, given the difficulty of predicting when and how many firm would develop a mature driverless technology, it makes little business sense to invest in testing capabilities for a currently nonexistent product.

[22]To understand the scale of these tests, the "Pre-Crash Typology" developed by the Department of Transportation, for example, consists of 66 crash scenarios testing how vehicles behave prior to crashes; General Motors developed an additional 44 scenarios testing crash behavior (Zhao (2016)).

Next, a quick study of the long and painful path to establishing modern experimentation standards in FDA clinical trials should temper any optimism about the prospects of properly enforcing information disclosure. The first federal requirement of trial registration did not come until 1997 and many types of clinical trials were not included in its purview until the passage of the FDA Amendments Act (FDAAA) in 2007. The process of expanding coverage to most types of trials was a slow one, with significant legislative updates happening as recently as 2016.

Regulators will likely have to contend with regulating automated vehicle manufacturers who have considerable leeway in selectively reporting those trips where a driver did not have to intervene with a corrective action.

Post-Market Oversight and Penalties A very frequent proposal for regulatory agencies is to strengthen their post-market oversight of product performance and regulate on the basis of real world impact. This is certainly a good thing, but the issue of course is that regulatory agencies frequently do not have the budget or firm cooperation to carry out adequate post-approval review of products.

This has been a problem for the FDA and caused a backlog of unreviewed cases. Indeed, a 2016 study by the Government Accountability Office showed a backlog of over 1,400 overdue post-market analyses (Kaplan, 2018). Weak post approval oversight has had its impact on compliance: only 13% of clinical trials followed the FDA Amendments Act's timely disclosure requirements following drug approval, and 38% of trials were not reported even 5 years after completion (Richardson, 2016). While there is supposed to a $10,000 daily fine for not reporting results within a year, Richardson (2016) writes that not once has a fine been levied. Finally, the most suggestive recent evidence of weak post-market oversight is found in the decades of regulatory failure in recognizing and stymieing the prescription of opioids. Heyward et al. (2020) show that opioid manufacturers collected non-representative data that left the FDA incapable of assessing whether a newly required safety program reduced the number of physicians making inappropriate prescriptions.

Even in the cases where regulators review product safety or efficacy after approval, they may face difficulty penalizing firms or pulling their technologies from the market. Courts rather than regulators control the outcomes from prosecuting delinquent firms, and the desired penalties are not always meted out. While the FDA can suggest drug companies to recall their products, it needs to sue them in court in cases where the companies refuse to comply. This is exceedingly rare given the regulatory body's limited resources.

There certainly exist regulatory bodies with richer sets of tools. But the ubiquity of the limitations discussed in this section motivates the focus on weak regulators, whose main levers are partial approval and costly standards of proof.

7 Concluding Remarks

I have explored the features of an approval process that arise when a regulator is concerned with avoiding mistakes and faces ambiguity over product quality about which a firm gives one-sided evidence. The combination of bureaucratic reputational concerns and ambiguity make regulators behave as if they minimize max regret. This gives a new explanation for why regulators may choose to drag their feet with approval decisions, how they tune their behavior to the evidence given, and why they may create regulatory sandboxes.

Here I studied the behavior of regulators who primarily control approval processes by choosing some form of partial approval. A natural continuation is to consider regulators with other effective levers like those described in Section 6.2. For example, an interesting extension is to allow the rate of information arrival to depend on the extent of approval awarded, which is possible if the regulator or third-parties can actively monitor post-approval product performance.

Another direction is to study how a social planner should influence a regulator's incentives. Such an extension is possible if the current model is enriched to allow the regulator's approval rule to influence firms' decisions to innovate.

Note that the regulator's concern with type II errors is that the true qualities of an already created innovation would eventually be revealed, despite her approval decision. In the current model, the regulator does not understand the innovation process; she entertains the possibility that despite her choice of approval rule, a product of any type can arise in practice. Still, if an onerous regulatory process were to dissuade a firm from innovating in the first place, the regulator would not be held accountable. Therefore if the regulator (were to believe she) could partially influence the innovation process, the interests of the regulator and a social planner may diverge. The regulator, for example, may try to stifle innovation to influence the distribution of innovating firms. A useful direction for future work is to consider how a planner would design a regulator's incentives, when the latter's choices affect the innovation process.

While the paper focuses on explaining the behavior of regulators, similar forces of ambiguity, screening and evidence manipulation may arise between a venture capitalist and startups seeking funding. The desire to minimize regret here may come from the fact that VCs find themselves in a competitive environment where a common pool of investor funds and deal flow depend on a venture capitalist's past performance. Adapting the current model to this setting may shed light on VC investment behavior and contract design.

A Appendix: Regret Minimization Representation with Stochastic Outcomes

A firm's type θ might not pin down its private benefit or the regulator's outcome. In the notation introduced in Section 2, $b(\theta)$ and $\tilde{v}(\theta)$ may be non-degenerate random-variables. Since the firm is risk neutral, its choice of evidence provision only depends on $E[b(\theta)]$. Here I show that a risk neutral regulator who is evaluated on the basis of consumer protection behaves *as if* she minimizes worst-case expected regret even when the outcomes to approval are stochastic. Define $v : \Theta \to \Delta\mathbb{R}$ such that

$$v(\theta) \equiv \begin{cases} \tilde{v}_0 + C - E[f(\tilde{v}_0 - \tilde{v}(\theta))|\tilde{v}(\theta) \leq \tilde{v}_0], \text{ if } \tilde{v}(\theta) \leq \tilde{v}_0 \\ \tilde{v}_0 + C + E[g(\tilde{v}(\theta) - \tilde{v}_0)|\tilde{v}(\theta) > \tilde{v}_0], \text{ if } \tilde{v}(\theta) > \tilde{v}_0 \end{cases}$$

where I suppose, without loss of generality, that $C > 0$ is a constant chosen so that $v \geq 0$. Let $v_0 \equiv \tilde{v}_0 + C$. Note that $v(\theta)$ is a binary random variable with $v(\theta) > v_0$ if and only if $\tilde{v}(\theta) > \tilde{v}_0$.

A regulator who minimizes expected worst case losses from type I and type II errors chooses $y \in [0, 1]$ to solve

$$\begin{aligned}
&\inf_{y\in[0,1]} \sup_{\theta\in\Theta} E[f(\tilde{v}_0 - \tilde{v}(\theta))y\mathbb{1}_{\tilde{v}_0 \geq \tilde{v}(\theta)} + g(\tilde{v}(\theta) - \tilde{v}_0)(1-y)\mathbb{1}_{\tilde{v}_0 < \tilde{v}(\theta)}] \\
&= \inf_{y\in[0,1]} \sup_{\theta\in\Theta} yE[f(\tilde{v}_0 - \tilde{v}(\theta))|\tilde{v}_0 \geq \tilde{v}(\theta)]Pr(\tilde{v}_0 \geq \tilde{v}(\theta)) + (1-y)E[g(\tilde{v}(\theta) - \tilde{v}_0)|\tilde{v}_0 < \tilde{v}(\theta)]Pr(\tilde{v}_0 < \tilde{v}(\theta)) \\
&= \inf_{y\in[0,1]} \sup_{\theta\in\Theta} yE[v_0 - v(\theta)|v_0 \geq v(\theta)]Pr(v_0 \geq v(\theta)) + (1-y)E[v(\theta) - v_0|v_0 < v(\theta)]Pr(v_0 < v(\theta)) \\
&= \inf_{y\in[0,1]} \sup_{\theta\in\Theta} E[(v_0 - v(\theta))y\mathbb{1}_{v_0 \geq v(\theta)} + (v(\theta) - v_0)(1-y)\mathbb{1}_{v_0 < v(\theta)}] \\
&= \inf_{y\in[0,1]} \sup_{\theta\in\Theta} E[\max\{v_0, v(\theta)\} - (1-y)v_0 - yv(\theta)].
\end{aligned}$$

B Appendix: Characterizing all Optimal Approval Rules in the Consumer Protection Problem

I state a generalized version of Proposition 1 that characterizes all optimal rules in the consumer protection problem.

Let $\overline{y}(\tau) \equiv \min\{1, \frac{R^C(y^*)}{v_0 - l(\tau)}\}$ and $\underline{y}(\tau) \equiv \max\{0, 1 - \frac{R^C(y^*)}{m(\tau) - v_0}\}$.[23]

Let $\mathcal{Y}$ be the set of all admissible, nondecreasing approval rules $\tau \to [0,1]$ that have the property that for all τ such that $l(\tau) < v_0 < m(\tau)$, $y(\tau) \in [\underline{y}(\tau), \overline{y}(\tau)]$. Note that $y^* \in \mathcal{Y}$.

PROPOSITION **6** (Proposition 1: Generalized)**.** *Suppose that either*

1. $l(\cdot), m(\cdot)$ *are nondecreasing, or*
2. $l(\cdot), m(\cdot)$ *are differentiable and* $m'(\tau)(v_0 - l(\tau)) + l'(\tau)(m(\tau) - v_0) \geq 0$ *for all* τ*.*

Then $\mathcal{Y}$ *is the set of approval rules which minimize regret among all undominated and implementable approval rules. In particular,* $y^* \in \mathcal{Y}$*.*

We can similarly characterize all optimal rules when y^* is not monotonic.

Fix some $R \in \mathbb{R}_+$. Let $\mathcal{Y}_R$ be the set of all admissible, nondecreasing approval rules $\tau \to [0,1]$ that have the property that for all τ such that $l(\tau) < v_0 < m(\tau)$, $y(\tau) \in [\underline{y}(\tau)_R, \overline{y}(\tau)_R]$, where $\overline{y}(\tau) = \min\{1, \frac{R}{v_0 - l(\tau)}\}$ and $\underline{y}(\tau) = \max\{0, 1 - \frac{R}{m(\tau) - v_0}\}$.

The next corollary (the proof of which closely follows the characterization of all optimal rules in case where y^* is monotonic) suggests a procedure for finding approximately optimal regret minimization schemes akin to the example shown in Figure 4: increment the regret level R until there exists a nondecreasing selection in the area between $\underline{y}(\cdot)_R$ and $\overline{y}(\cdot)_R$.

COROLLARY **1.** $R \in \mathbb{R}_+$ *is the minimum regret achievable by an admissible, implementable approval rule if for all* $R' < R$, $\mathcal{Y}_{R'}$ *is empty but* $\mathcal{Y}_R$ *is nonempty. A regret minimizing approval rule exists.*

[23] Notice that for every τ such that $l(\tau) < v_0 < m(\tau)$, $\overline{y}(\tau) \geq \underline{y}(\tau)$: By Lemma 1, $R^C_\tau(y^*) = \frac{(m(\tau) - v_0)(v_0 - l(\tau))}{m(\tau) - l(\tau)}$, so $R^C(y^*) \geq \frac{(m(\tau) - v_0)(v_0 - l(\tau))}{m(\tau) - l(\tau)} \implies R^C(y^*)(m(\tau) - v_0) + R^C(y^*)(v_0 - l(\tau)) \geq (m(\tau) - v_0)(v_0 - l(\tau)) \implies \frac{R^C(y^*)}{v_0 - l(\tau)} \geq 1 - \frac{R^C(y^*)}{m(\tau) - v_0}$.

C Appendix: Limited Manipulation

While there are instances of egregious manipulations of evidence in the past, outright nondisclosure of failed experiments are less likely to occur in the present-day FDA's more formal clinical trial approval process. Still, pharmaceutical companies may still have some limited room to costlessly inflate the realized efficacy and safety metrics in trials. For example, they may halt trials prematurely the moment positive results are recorded. In this section, I consider a model where the firm cannot hide negative trials, but can design experiments in such a way as to artificially boost its in-trial success rate. I solve for an optimal approval rule in a simple example.

In this model, letting $0 < \epsilon < 1$, the firm's type is effectively three dimensional: $(q, k, \psi) \in \Theta \equiv [\underline{q}, \overline{q}] \times [\underline{k}, \overline{k}] \times [0, \epsilon]$. The efficacy of the drug in question is $q \in [\underline{q}, \overline{q}] \subset (0, 1)$ (e.g., cures an illness with probability q) and the cost of running a trial is $k \in [\underline{k}, \overline{k}]$. The new term $\psi \in [0, \epsilon]$ captures the firm's capacity for manipulation. In particular, the firm of type (q, k, ψ) can design a trial so that it returns a success with probability $q' \in [0, q + \psi]$.

It is convenient to collapse the experimentation and reporting stage of the firm's action and modify its action space as follows. Assume that a firm with type (q, k, ψ) has an action space given by $[0, q + \psi] \times \mathbb{R}_+$, where the first dimension represents the success rate of firm's trials and the second dimension is the number of trials the firm runs.[24] The regulator commits to a decision rule $t : [0, 1] \times \mathbb{R}_+ \to \mathbb{R}_+$, mapping the firm's action to a date of approval. In my example here, the regulator's benefit of approving a type (q, k, ψ) only depends on q, with $v(q, k, \psi) = q$ for some nondecreasing f. If the regulator does not approve the firm, her payoff is v_0. With an experiment of size n and a test of efficacy $q' \in [0, q + \psi]$, the firm gets a payoff of $e^{-t(q',n)}b(q, q') - kn$, where $b(q, q')$ is the firm's benefit of immediate approval. In this example, the firm's benefits do not depend on its costs of trial k.

Note that the firm can choose actions of the kind $(q', 0) \in [0, q + \psi] \times \mathbb{R}_+$: the firm can display its type without producing a positive mass of trials. Under the small costs interpretation, this means that the total cost of producing the number of trials to confirm the firm's chosen success rate with high confidence is negligible, at least compared to the benefits of approval.

In particular, if a firm had a drug of efficacy $q > q_0 + \epsilon$, then it can "costlessly"

[24] This is meant to capture the fact that the firm can run small trials at little cost, which side steps the complication that realized outcomes may differ from the firm's chosen success rate in small samples. A game closer in description to the one analyzed earlier would be as follows. While the second period (experimentation stage) of the earlier game would be the same as before, one could modify the firm's action space in the third period (reporting stage). Suppose that a firm that has generated x_0 failures and x_1 successes must report $(x_0, x_1) \in \mathbb{R}^2_+$. But to easily show existence of equilibria, it is also be helpful to assume that the firm can run an infinitesimal experiment at no cost to reveal its chosen success rate $q' \in [0, q + \psi]$. The firm could then run very small trials in sequence to approximately reach its desired success rate.

prove to the regulator that its drug is superior to the existing standard. Even if it was manipulating evidence, its type would have to be greater than q_0 to produce an experiment with success rate q.

Unlike the cases addressed in the previous section, the firm now has a two dimensional action space which changes the nature of the regulator's design problem. Let $h \equiv q + \psi$ be the firm's *largest possible experimental success rate (LPESR)*.

Case 1: $b(q, q') = q$**.** In this case, the private benefit of approval of the firm depends only on the true efficacy of its drug.

Note that any two types with the same LPESR have the same action space.

By Assumption 1, we can restrict search for regret minimizing mechanisms to IC approval rules which give the same delay to all firm types with the same effective type, τ, and LPESR, h (call the set of such types $\Theta_{\tau,h}$).

Only types with efficacy in the range $[q - \epsilon, q]$ can have an LPESR of q. Since the bounds of this interval are increasing in q, for any given τ, $l(\tau, h) \equiv \inf_{\theta \in \Theta_{\tau,h}} v(\theta)$ and $m(\tau, h) \equiv \sup_{\theta \in \Theta_{\tau,h}} v(\theta)$ are nondecreasing in h (when $\Theta_{\tau,h}$ is non-empty and non-singleton). Therefore the pointwise regret minimizing approval rule $y^*(\tau, h)$ is nondecreasing in both arguments. To extend the definition of y^* on the range where $\Theta_{\tau,h} = \emptyset$, let $y^*(\tau, h) = 0$ if $h - \epsilon > \tau\overline{k}$ and $y^*(\tau, h) = 0$ if $h < \tau\underline{k}$. If y^* is implementable, then it is clearly the optimal solution.

Let $n^*(\tau, h) \equiv y^*(\tau, h)\tau - \int_{\underline{\tau}}^{\tau} y^*(s, h)ds$ for all $\tau \in \mathcal{T}$ and $h \in [\underline{q}, \overline{q}+\epsilon]$. Let $\theta \equiv (q, k, \psi)$ and $\tilde{\theta} \equiv (\tilde{q}, \tilde{k}, \tilde{\psi})$. Let $\tau \equiv \frac{q}{k}$ and $\tilde{\tau} \equiv \frac{\tilde{q}}{\tilde{k}}$, and $h \equiv q + \psi$ and $\tilde{h} \equiv \tilde{q} + \tilde{\psi}$. Supposing $\tilde{h} \leq h$, to prove that n^* implements y^*, it suffices to show

$$y^*(\tau, h)\tau - n^*(\tau, h) \geq y^*(\tilde{\tau}, \tilde{h})\tau - n^*(\tilde{\tau}, \tilde{h}).$$

This is true if and only if

$$\int_{\underline{\tau}}^{\tau} y^*(s, h)ds \geq y^*(\tilde{\tau}, \tilde{h})(\tau - \tilde{\tau}) + \int_{\underline{\tau}}^{\tilde{\tau}} y^*(s, \tilde{h})ds$$

$$\iff \int_{\underline{\tau}}^{\tau} y^*(s, h) - y^*(s, \tilde{h})ds \geq y^*(\tilde{\tau}, \tilde{h})(\tau - \tilde{\tau}) + \int_{\tau}^{\tilde{\tau}} y^*(s, \tilde{h})ds$$

Since y^* is nondecreasing in both arguments, the left side is always non-negative and the right side is always non-positive. Therefore y^* is the optimal implementable mechanism.

Reverting back to the original action space, where the firm chooses a level of manipulation and size of experiment, the firm always inflates its experimental success rate as much as possible under the regulator's optimal approval rule.

Case 2: $b(q, q') = g(q')$, where g is some nondecreasing function. In this case, the private benefit of the firm is only a function of its perceived efficacy in clinical trials. Note that a special case is that g is constant, so the private value of approval is same for all types. This would happen in a regime with little market regulation, where the company's ability to promote its drug is independent of its true or perceived efficacy.

In this case, note that $\Theta_h = \{(q, k, \psi) | q + \psi = h\}$ consists of types with identical action spaces and preferences over actions and outcomes. This means that $l(\tau, h) = h - \epsilon$ and $m(\tau, h) = h$, i.e., $y^*(\cdot, h)$ does not depend on its first argument.

Letting q' be the firm's chosen experimental success rate, this implies that an optimal mechanism is as follows: approve the firm immediately if $q' - \epsilon > v_0$, never allocate if $q' < v_0$, and otherwise delay the firm so that they are left with a fraction $\frac{q'-v_0}{\epsilon}$ of their private values.[25] Moreover, since the allocation depends entirely on the firm's experimental success rate, there is no screening through experiment size, so the standard of proof that implements the optimal mechanism is $n = 0$. The firm's optimal response is to always inflate their experimental success rate to their LPESR.

[25]The arguments are unchanged if the regulator's payoff to approving a firm of type q is $f(q)$, where f is increasing and non-negative

D Appendix: Example of Non-Monotonicities in Consumer Protection Problem

Non-monotonicities can naturally arise in the regulator's problem when the regulator's interests and consumer interests are divergent. Here I consider an example where the regulator is evaluated purely based on the average quality of products used in the market, whereas consumers care about both quality and price.

For illustration, recall the example in Section 4.1 where the regulator knows only that the firm's success rate q falls in an interval $[\underline{q}, \overline{q}]$ and the cost of trials $k \in [\underline{k}, \overline{k}]$. Where I assumed before that the firm's benefit of approval is identical across types, I now allow benefits to vary.

In particular, suppose there is a unit mass of consumers in the market who vary by how much they value a product's efficacy. If a firm prices its product at p, a consumer of type $\omega \in \Omega$ values the product at $\omega \cdot q - p$. There is an existing product with success rate q_0 produced in a perfectly competitive market at price p_0.

The firm enters this market as a monopolist. I assume that it has no fixed costs of production and that its marginal cost is precisely its cost of running trials, k. Upon approval, the firm chooses its price $p^*(q, k)$ optimally, given its product quality q and cost of production k. A consumer of type ω will purchase the firm's product only if $\omega \cdot q - p \geq \omega \cdot q_0 - p_0$. By the assumption of perfect competition in the existing market, the price of the existing product will remain at p_0 (industry marginal cost) regardless of the monopolist's behavior. It is easy to see that the firm's demand for its product, $D(p, q)$ is nondecreasing in q and $D(p^*(q, k), q)$ is nonincreasing in k.

Suppose that the regulator's payoff is the expected quality of consumed products, $D(p^*(q, k), q)q + (1 - D(p^*(q, k), q))q_0$. For example, the regulator might only care about minimizing the expected accident rates while drivers themselves are willing to trade-off safety for low prices.

Consider different product safety levels $q' < q'' < q_0$. While a firm with success rate q'' may seem to be preferable to a firm with success rate q', it may actually be less desirable for the regulator to approve if $D(p^*(q'', k), q'')$ is sufficiently larger than $D(p^*(q', k), q')$. The intuition is the same as in the imperfect markets case: an inferior product with greater market penetration may end up doing more damage on the whole than a low quality product.

Indeed, when the firm can hide evidence, Figure D gives a numerical example where $l(\cdot)$ is decreasing in some range of effective types.

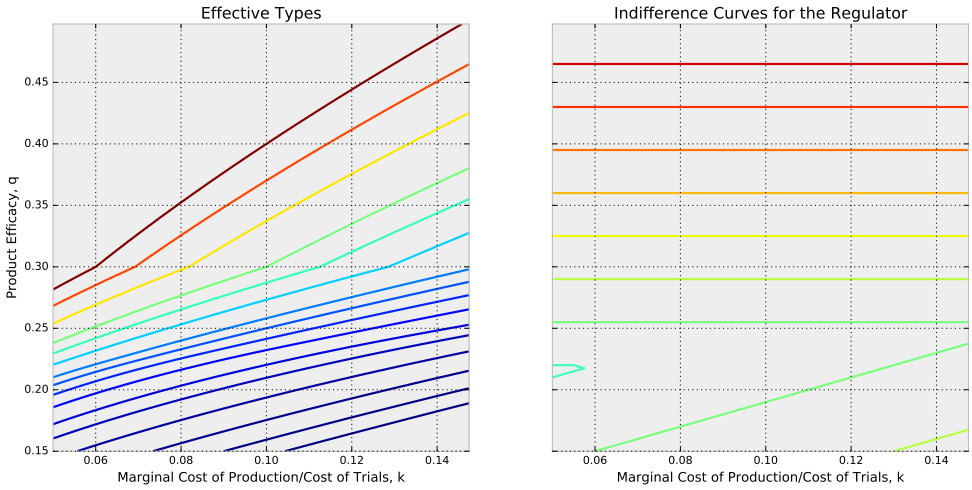

Figure 6: Following the 'Misaligned Preferences' example, these graphs depict effective types and regulator preferences when consumer types, $\omega \in \Omega$, are uniformly distributed on $[1, 2]$; $[\underline{k}, \overline{k}] = [0.05, 0.15]$; and $p_0 = q_0 = 0.3$. Types along the contour lines in the left plot have the same effective type. The contour lines in the plot to the right are the regulator's iso-utility curves over the firm's type space. In both plots, redder lines represent an increase in value and bluer lines a decrease. Note that as effective types are increasing, the regulator's worst case outcome per effective type is initially decreasing.

E Appendix: Omitted Proofs

Proof of Lemma 2. Fix an $\tilde{n}$ such that $(\tilde{x}, \tilde{n})$ is an incentive compatible menu. IC implies that,

$$\tau(\hat{\theta})\tilde{x}(\hat{\theta}) - \tilde{n}(\hat{\theta}) \geq \tau(\hat{\theta})\tilde{x}(\theta) - \tilde{n}(\theta)$$
$$\tau(\theta)\tilde{x}(\theta) - \tilde{n}(\theta) \geq \tau(\theta)\tilde{x}(\hat{\theta}) - \tilde{n}(\hat{\theta}).$$

Subtracting the first inequality from the second,

$$(\tau(\theta) - \tau(\hat{\theta}))\tilde{x}(\theta) \geq (\tau(\theta) - \tau(\hat{\theta}))\tilde{x}(\hat{\theta}),$$

so $\tilde{x}(\theta) \geq \tilde{x}(\hat{\theta})$. □

Proof of Lemma 3. One direction comes from Lemma 2. For the other direction, suppose y is nondecreasing. Let $n(\tau) = \tau y(\tau) - \underline{\tau} y(\underline{\tau}) - \int_{\underline{\tau}}^{\tau} y(t)dt$. I claim that this choice of n implements y. Fix $\tau, \tau' \in \mathcal{T}$. Since y is nondecreasing,

$$y(\tau')(\tau' - \tau) \geq \int_{\tau}^{\tau'} y(t)dt$$
$$\implies y(\tau')\tau' - y(\tau)\tau \geq \int_{\tau}^{\tau'} y(t)dt + y(\tau')\tau - y(\tau)\tau$$
$$\implies \tau(y(\tau) - y(\tau')) \geq n(\tau) - n(\tau'),$$

which is the IC condition after multiplying through by $c(\theta)$ for any θ such that $\tau(\theta) = \tau$. □

Proof of Lemma 4. It is straightforward to check that if y is undominated, then it is admissible. For the other direction, suppose y is admissible, and take any other approval rule $\hat{y}$.

If $l(\tau) = m(\tau) = v_0$, then for every $\theta \in \Theta_\tau$, $y(\theta)v(\theta) + (1 - y(\theta))v_0 = \hat{y}(\theta)v(\theta) + (1 - \hat{y}(\theta))v_0$.

If $l(\tau) \geq v_0$ and $m(\tau) > v_0$, then $v(\theta) \geq v_0$ for all $\theta \in \Theta_\tau$, so $y(\theta)v(\theta) + (1 - y(\theta))v_0 = v(\theta) \geq \hat{y}(\theta)v(\theta) + (1 - \hat{y}(\theta))v_0$. The case where $m(\tau) \leq v_0$ and $l(\tau) < v_0$ is handled similarly.

If $l(\tau) < v_0 < m(\tau)$ and $y(\tau) > \hat{y}(\tau)$, then there exists some $\theta \in \Theta_\tau$ such that $v(\theta) > v_0$, so $y(\theta)v(\theta) + (1 - y(\theta))v_0 > \hat{y}(\theta)v(\theta) + (1 - \hat{y}(\theta))v_0$. And if $y(\tau) < \hat{y}(\tau)$, there exists some $\theta \in \Theta_\tau$ such that $v(\theta) < v_0$, so again $y(\theta)v(\theta) + (1 - y(\theta))v_0 > \hat{y}(\theta)v(\theta) + (1 - \hat{y}(\theta))v_0$.

Therefore no $\hat{y}$ dominates y. □

Proof of Proposition 6. Conditions 1 and 2, along with Lemma 1 ensure that y^* is nondecreasing. Then by Lemma 3, y^* is implementable. By Lemma 4, y^* is undominated. Therefore $y^* \in \mathcal{Y}$.

If $y \in \mathcal{Y}$, then

$$\begin{aligned} R^C_\tau(y(\tau)) &= \max\{y(\tau)(v_0 - l(\tau)), (1 - y(\tau))(m(\tau) - v_0)\} \\ &\leq \max\{\overline{y}(\tau)(v_0 - l(\tau)), (1 - \underline{y}(\tau))(m(\tau) - v_0)\} \\ &\leq R^C(y^*). \end{aligned}$$

Since y^* minimizes $R^C_\tau(\cdot)$ for every τ, it also minimizes $R^C(\cdot)$ in the space of admissible approval rules. The same is then true for y.

If $y \notin \mathcal{Y}$ is undominated and nondecreasing, then there is a $\tau \in \mathcal{T}$ such that $R^C_\tau(y(\tau)) > R^C(y^*)$. □

Proof of Proposition 3. The proof proceeds in three steps.

Step 1: Approval rules can be improved to a CRCA rule with less regret.

First note that,

$$
\begin{aligned}
&\sup_{\tau} \max_{\theta\in\Theta_\tau} \max\{v_0, v(\theta)\} - [y(\tau)v(\theta) + (1-y(\tau))v_0)] + c(\theta)n(\tau) \\
=&\sup_{\tau} \max_{\substack{\theta\in\Theta_\tau \\ s\in[-d(\tau),D(\tau)]}} \max\{v_0, b(\theta)+s\} - [y(\tau)(b(\theta)+s) + (1-y(\tau))v_0)] + c(\theta)n(\tau) \\
=&\sup_{\tau} \max_{\substack{c\in[\underline{c},\overline{c}] \\ s\in[-d(\tau),D(\tau)]}} \max\{y(\tau)(v_0 - c\tau - s), (1-y(\tau))(c\tau + s - v_0)\} + cn(\tau) \\
=&\sup_{\tau} \max_{\substack{c\in[\underline{c},\overline{c}] \\ s\in[-d(\tau),D(\tau)]}} \max\{y(\tau)(v_0 - c\tau - s), (1-y(\tau))(c\tau + s - v_0)\} + c(y(\tau)\tau - y(\underline{\tau})\underline{\tau} - \int_{\underline{\tau}}^{\tau} y(t)dt) \\
=&\sup_{\tau} \max_{\substack{c\in[\underline{c},\overline{c}] \\ s\in[-d(\tau),D(\tau)]}} \max\{y(\tau)(v_0 - s), c\tau + (1-y(\tau))(s - v_0)\} - c(y(\underline{\tau})\underline{\tau} + \int_{\underline{\tau}}^{\tau} y(t)dt) \\
=&\sup_{\tau} \max_{c\in[\underline{c},\overline{c}]} \max\{y(\tau)(v_0 + d(\tau)), c\tau + (1-y(\tau))(D(\tau) - v_0)\} - c(y(\underline{\tau})\underline{\tau} + \int_{\underline{\tau}}^{\tau} y(t)dt) \\
=&\sup_{\tau} \max \begin{cases} y(\tau)(v_0 + d(\tau)) - \underline{c}(y(\underline{\tau})\underline{\tau} + \int_{\underline{\tau}}^{\tau} y(t)dt) \\ (1-y(\tau))(D(\tau) - v_0) + \overline{c}(\tau - y(\underline{\tau})\underline{\tau} - \int_{\underline{\tau}}^{\tau} y(t)dt) \end{cases} \\
\equiv&\sup_{\tau} \max\{R_{\tau,v_0}(y), R_{\tau,v_1}(y)\}
\end{aligned}
$$

Let $R \equiv R(y)$. Consider the largest p such that $p(v_0 + d(\underline{\tau}) - \underline{c}\underline{\tau}) \leq R$. Such a p exists by assumption that $R(y) = R$ (e.g., $p \geq y(\underline{\tau})$).

Case 1: If $p \geq y(\overline{\tau})$, consider $y' = y(\overline{\tau})$.

For $\tau \in \mathcal{T}$ where $R_{\tau,v_0}(y') \geq R_{\tau,v_1}(y')$, $R_\tau(y') \leq R_{\underline{\tau}}(y') \leq R$. The first inequality is because y' is constant and $R_{\tau,v_0}(y')$ is decreasing in τ; the second inequality is because $y(\underline{\tau}) \leq y(\overline{\tau}) \leq p$, and $R_{\underline{\tau}}(y) \leq R$ and $R_{\underline{\tau}}(p) \leq R$, and the range of allocation probabilities for which $R_{\underline{\tau}}(y) \leq R$ is convex.

For τ where $R_{\tau,v_0}(y') \leq R_{\tau,v_1}(y')$, $R_\tau(y') \leq R_{\overline{\tau}}(y') \leq R_{\overline{\tau}}(y) \leq R$. The first inequality is because, first, y' is constant and $R_{t,v_1}(y')$ is nondecreasing in t since $\overline{c}$ and $\overline{v}$ are; and moreover, $R_{t,v_1}(y') \geq R_{t,v_0}(y')$ for all $t > \tau$, since the latter is decreasing in τ. The second inequality follows from the fact that $n'(\overline{\tau}) = 0 \leq n(\overline{\tau})$ while $y'(\overline{\tau}) = y(\overline{\tau})$.

Case 2: If $p < y(\overline{\tau})$, let $\tilde{y}$ be the solution to the integral equation $R = R_{\tau,v_0}(y)$, with $\tilde{y}(\underline{\tau}) = p$.

Note that it is easy to verify that $\tilde{y}$ is continuous, and rearranging the terms shows that $\tilde{y}$ is a function of differentiable terms and is therefore differentiable. Next, Picard's theorem says that the solution to the differential equation with the given boundary condition is unique, so the solution to the initial integral equation is also unique. Finally, its easy to see that $\tilde{y}$ is strictly increasing by rearranging the differential equation and seeing the derivative of $\tilde{y}$ is positive everywhere. There are two cases to consider.

Subcase 1: $\tilde{y}(\tau') = y(\overline{\tau})$ at some τ' such that on $[\underline{\tau}, t'], R_{v_0,\tau}(\tilde{y}) \geq R_{v_1,\tau}(\tilde{y})$. Then setting $y' = \tilde{y}$ on $[\underline{\tau}, \tau']$ and $y' = y(\overline{\tau})$ on $[\tau', \overline{\tau}]$ satisfies the properties required by the proposition (verifying that $R(y') = R$ is by same arguments as before).

Subcase 2: There is a τ' and $d > 0$ such that 1) $R_{v_0,\tau}(\tilde{y}) \geq R_{v_1,\tau}(\tilde{y})$ on $[\underline{\tau}, \tau']$, 2) $R_{v_0,\tau}(\tilde{y}) < R_{v_1,\tau}(\tilde{y})$ for $\tau \in (\tau', \tau' + d]$, 3) $\tilde{y}(\tau') < y(\overline{\tau})$.

Consider any nondecreasing extension of $\tilde{y}$, $\hat{y} : [\underline{\tau}, \tau' + \epsilon) \to [0, 1]$. That is to say, $\hat{y} = \tilde{y}$ on $[\underline{\tau}, \tau']$ and is nondecreasing over its full domain. I want to show that $R_\tau(\hat{y}) > R$ at some $\tau \in (\tau', \tau' + \epsilon)$.

If $\hat{y}$ is discontinuous at τ', it is clear that $\lim_{\tau \to \tau'_+} R_{v_0,\tau}(\hat{y}) > R$ so $R_\tau(\hat{y}) > R$ for τ sufficiently close to τ'. Suppose then that $\hat{y}$ is continuous at τ'.

If $\hat{y} = \tilde{y}$ on $(\tau', \tau' + \epsilon)$ for some $\epsilon > 0$, by assumption of Subcase 2, $R_\tau(\tilde{y}) = R_{v_1,\tau}(\tilde{y}) > R$ for some τ in this region.

Suppose now that $\hat{y} > \tilde{y}$ on some $(\tau', \tau' + \epsilon)$. If $R_\tau(\hat{y}) < R$ on this region, then of course $R_{\tau,v_0}(\hat{y}) < R$ as well. This implies that there exists another $\hat{y}'$ such that $\hat{y}' = y'$ on $[\underline{\tau}, \tau']$ and $\hat{y}' \geq \hat{y}$ on $(\tau', \tau + \epsilon]$, with $R_{\tau,v_0}(\hat{y}') = R$ on $[\underline{\tau}, \tau' + \epsilon]$. But this is a contradiction since $\hat{y}' \neq \tilde{y}$, and I showed that the latter is the unique solution to the integral equation $R = R_{\tau,v_0}$ on this region.

Next suppose that $\hat{y} < \tilde{y}$ and that $D(\tau) \geq v_0$ on an interval $[\tau', \tau' + \epsilon)$. Then $R_{v_1,\tau}(\hat{y}) \geq R_{v_1,\tau}(\tilde{y}) > R_{v_0,\tau}(\tilde{y}) = R$ on $[\tau', \tau' + \epsilon)$.

Finally, suppose that $\hat{y} < \tilde{y}$ and that $D(\tau) < v_0$ on an interval $(\tau', \tau' + \epsilon]$. Let $\hat{y}(\tau' + \epsilon) = \hat{y}(\tau') + x$, for some $x \geq 0$.

$$
\begin{aligned}
&R_{\tau'+\epsilon}(\hat{y}) \geq R_{\tau'+\epsilon,v_1}(\hat{y}) \\
&= (1 - \hat{y}(\tau' + \epsilon))(D(\tau' + \epsilon) - v_0) + \overline{c}(\tau' + \epsilon - \hat{y}(\underline{\tau})\underline{\tau} - \int_{\underline{\tau}}^{\tau'+\epsilon} \hat{y}(t)dt) \\
&\geq (1 - \hat{y}(\tau') - x)(D(\tau' + \epsilon) - v_0) + \overline{c}(\tau' + \epsilon - \hat{y}(\underline{\tau})\underline{\tau} - \int_{\underline{\tau}}^{\tau'} \hat{y}(t)dt - \int_{\tau'}^{\tau'+\epsilon} (\hat{y}(\tau') + x)dt) \\
&= (1 - \hat{y}(\tau') - x)(D(\tau' + \epsilon) - v_0) + \overline{c}(\tau' + \epsilon - \hat{y}(\underline{\tau})\underline{\tau} - \int_{\underline{\tau}}^{\tau'} \hat{y}(t)dt - \epsilon(\hat{y}(\tau') + x)) \\
&= [(1 - \hat{y}(\tau'))(D(\tau' + \epsilon) - v_0) + \overline{c}(\tau' - \hat{y}(\underline{\tau})\underline{\tau} - \int_{\underline{\tau}}^{\tau'} \hat{y}(t)dt)] - x(D(\tau' + \epsilon) - v_0) + \overline{c}\epsilon(1 - \hat{y}(\tau') - x) \\
&\geq [(1 - \hat{y}(\tau'))(D(\tau') - v_0) + \overline{c}(\tau' - \hat{y}(\underline{\tau})\underline{\tau} - \int_{\underline{\tau}}^{\tau'} \hat{y}(t)dt)] - x(D(\tau' + \epsilon) - v_0) + \overline{c}\epsilon(1 - \hat{y}(\tau') - x) \\
&= R - x(D(\tau' + \epsilon) - v_0) + \overline{c}\epsilon(1 - \hat{y}(\tau') - x) \\
&> R
\end{aligned}
$$

Now let $y' = \tilde{y}$ on $[\underline{\tau}, \tau']$ and $y' = y'(\tau')$ on $(\tau', \overline{\tau}]$, and let $\tau^* = \inf\{\tau \in \mathcal{T} | y(\tau) > y'(\tau)\}$.

I first consider the case where $\tau^* \in [\underline{\tau}, \tau']$. Suppose first that $y(\tau) = y'(\tau)$ for all $\tau < \tau^*$. If y jumped discontinuously at τ', $R_{\underline{\tau},v_0}(y)$ would jump at τ from R, contradicting $R(y) = R$. If y is continuous at τ, I would get the contradiction that $R_\tau(y) > R$ for some τ close to τ^*, by construction of y' (appealing to the fact that I would then be able to construct another $y'' \neq y'$ that also solves $R_{v_0,\tau}(y') = R$, contradicting the uniqueness of the solution). Suppose now that $y(\tau) < y'(\tau)$ for some $\tau < \tau^*$. Since y is increasing

and y' is continuous, this implies that $y < y$ in some interval $(\tau - \epsilon, \tau)$. But then $\lim_{\tau \to \tau_+^*} R_\tau(y) \geq \lim_{\tau \to \tau_+^*} R_{\tau^*, v_0}(y) > R_{\tau^*, v_0}(y') = R$, another contradiction.

Next, suppose $\tau^* \in (\underline{\tau}, \tau')$. If $y = y'$ for all $\tau < \tau*$, I know that since every extension of $\tilde{y}$ past τ' has regret greater than R at some point, this would contradict $R(y) = R$. But the case where $y(\tau) < y'(\tau)$ for some $\tau < \tau^*$ gives the same contradiction as before.

I have shown that subcase 2 is an impossibility, given the assumptions of case 2, proving step 1.

Step 2: Existence of Optimal Rule. First I restrict search for optimal approval rules to the space of continuous approval rules. Since the adversary's action space is compact, for every continuous rule y, the adversary has an optimal choice of $\theta \in \Theta$.

Now $R_{\tau, v_0}(y)$ and $R_{\tau, v_1}(y)$ are both continuous in τ and y (under the infinity norm metric), so $\max_{\tau \in \mathcal{T}} R_{\tau, v_0}(y)$ and $\max_{\tau \in \mathcal{T}} R_{\tau, v_1}(y)$ are continuous in y. Since the space of continuous approval rules is compact under the infinity norm, the regulator's problem restricted to this space admits a solution.

But by the earlier proposition, for any y, there is a $y' \in \overline{\mathcal{Y}}$ with $R(y') = R(y)$, and since y' is continuous, the earlier restriction was without loss of generality.

Step 3: Uniqueness of Optimal Rule Let $\tilde{y}$ be the solution to $R = R_{\tau, v_0}(y)$. Note that $\frac{d}{d\tau} R_{\tau, v_1}(\tilde{y}) = -(\frac{d}{d\tau}\tilde{y})(D(\tau) - v_0) + (1 - \tilde{y}(\tau))\frac{d}{d\tau} D(\tau) + \overline{c}(1 - \tilde{y}(\tau)) \geq 0$. On the other hand, if y is a constant approval rule, $\frac{d}{d\tau} R_{\tau, v_1}(y) = (1 - y(\tau))\frac{d}{d\tau} D(\tau) + \overline{c}(1 - y(\tau)) \geq 0$ (with equality only if $y = 1$). Notice moreover that $R_{\tau, v_0}(y)$ for any $y \in \overline{\mathcal{Y}}$ is initially constant and then decreasing. Therefore, for any $y \in \overline{\mathcal{Y}}$, $R_{\tau, v_0}(y)$ and $R_{\tau, v_1}(y)$ cross only once, if at all.

Now let R be the value of the regulator's problem under an optimal approval rule. Take some $y \in \overline{\mathcal{Y}}$ with $R(y) = R$ and with τ' being the cutoff effective type at which y becomes constant. Suppose that the $R_{\overline{\tau}}(y) < R$. Then I can consider another approval rule $y' \in \overline{\mathcal{Y}}$ with $R_{\underline{\tau}, v_0}(y') = R - \epsilon$, which has the same cutoff effective type τ' as y. If ϵ is small enough, by continuity and the fact that $R_\tau(y) < R$ at any τ at which $R_{\tau, v_0}(y) = R_{\tau, v_1}(y)$, $R_{\overline{\tau}, v_1}(y') < R$, so $R(y') < R$, contradicting the fact that y was assumed to be optimal.

Therefore, $R(y) = R_{\overline{\tau}}(y)$. Since $R_{\tau, v_0}(y)$ is decreasing on the portion of y where it is flat, the preceding also implies $R_{\overline{\tau}}(y) = R_{\overline{\tau}, v_1}(y)$.

Now suppose y satisfies this property and has its cutoff point at $\tau' < \overline{\tau}$. If $v_0 - D(\overline{\tau}) - \overline{c}(\overline{\tau} - \tau') < 0$, then an approval rule $\hat{y}$ with regret $R(y)$ on $[\underline{\tau}, \tau']$ but cutoff at $\tau' + \epsilon$ for small ϵ will have the property that $R_{\overline{\tau}}(\hat{y}) < R$, which by the preceding paragraph, contradicts the optimality of y. [26] Similarly, if $v_0 - D(\overline{\tau}) - \overline{c}(\overline{\tau} - \tau') < 0$, an approval rule $\hat{y}$ with the same regret on $[\underline{\tau}, \tau']$ but cutoff at $\tau' - \epsilon$ for small ϵ will produce the same contradiction.

[26] The regret at $\overline{\tau}$ of the constant regret, constant allocation policy with cutoff at τ' is $(1 - \tilde{y}(\tau'))(D(\overline{\tau}) - v_0) + \overline{c}(\overline{\tau} - y(\underline{\tau})\underline{\tau} - \int_{\underline{\tau}}^{\tau'} \tilde{y}(t)dt - \tilde{y}(\tau')(\overline{\tau} - \tau'))$. Taking the derivative of this expression with respect to the cutoff τ' gives $\frac{d}{d\tau'}\tilde{y(\tau)}(v_0 - D(\overline{\tau}) - \overline{c}(\overline{\tau} - \tau'))$.

Let τ' solve $v_0 - D(\overline{\tau}) - \overline{c}(\overline{\tau} - \tau') = 0$, and let $y^* \in \overline{\mathcal{Y}}$ be the approval rule with minimum regret among those with cutoff τ. By the preceding arguments, this is the only rule which cannot be improved (within $\overline{\mathcal{Y}}$, and therefore in the space of all rules), and since an optimal rule exists, it must be y^*. □

Proof of Proposition 4. First consider the case where $h = 0$, i.e., the case where the regulator is judged only on the basis of consumer protection.

Let y be a solution to the regulator's problem. Suppose $y(\tau) \in (0,1)$ for $\tau \in (c,d)$, and $y = 0$ for $\tau < c$, and $y = 1$ for $\tau > d$.

Take a finite partition, discrete approximation y_d to y such that $V_B(y_d)$ is within ϵ of $V_B(y)$. Note this is possible since there are no atoms and y, being increasing, is continuous almost everywhere. Let $(p_0, \ldots, p_n)$ be the partition P such that y_d is constant on the sub-intervals (p_i, p_{i+1}), taking on the value I denote by z_i, with $p_0 = c$ and $p_n = d$.

Consider the following construction:

1. If y_d is already a cutoff policy, that gives the desired result.
2. Otherwise, consider the interval (p_{n-1}, p_n).

   If $\int_{p_{n-1}}^{p_n} (E[v(\theta)|\tau] - v_0) y_d(\frac{1}{2}(p_{n-1} + p_n)) \geq 0$, let $y_d^1 = 1$ on (p_{n-1}, p_n) and $y_d^1 = y_d$ outside of this interval.

   If $\int_{p_{n-1}}^{p_n} (E[v(\theta)|\tau] - v_0) y_d(\frac{1}{2}(p_{n-1} + p_n)) < 0$, let $y_d^1 = z_{n-2}$ on (p_{n-1}, p_n) and $y_d^1 = y_d$ outside of this interval.
3. Define a new partition $P^1 = (p_0^1, \ldots, p_{n-1}^1)$, where $p_{n-1}^1 = d$, $p_{n-2}^1 = p_{n-1}$, and $p_j^1 = p_j$ for all other j. That is, P^1 is the same as P, but with (p_{n-1}, p_n) and (p_{n-2}, p_{n-1}) merged.
4. Relabel y_d^1 as y_d and P' as P

It is clear that iterating these steps, with y_d^1 and P^1 in place of y_d and P respectively, one arrives at a a cutoff policy. Moreover, $V_B(y_d^1) \geq V_B(y_d)$ by construction, so V_B evaluated at the terminal cutoff policy is greater than $V_B(y_d)$. The latter was within ϵ of $V_B(y)$ and the choice of ϵ was arbitrary, completing the proof for the case of consumer protection.

Returning to the case where the regulator is judged on consumer protection and firm costs, note:

$$\int_{\underline{\tau}}^{\overline{\tau}} n(\tau) E[h(\theta)|\tau] dF_\tau(\tau) = \int_{\underline{\tau}}^{\overline{\tau}} (y(\tau)\tau + n(\underline{\tau}) - \underline{\tau} y(\underline{\tau}) - \int_{\underline{\tau}}^{\tau} y(t)dt) E[h(\theta)|\tau] dF_\tau(\tau),$$

so

$$\begin{aligned} V_B(y,n) = &\int_{\underline{\tau}}^{\overline{\tau}} y(\tau)(E[v(\theta)|\tau] - v_0) - (y(\tau)\tau + n(\underline{\tau}) - \underline{\tau} y(\underline{\tau}))E[h(\theta)|\tau])dF_\tau(\tau) \\ &+ \int_{\underline{\tau}}^{\overline{\tau}} \int_{\underline{\tau}}^{\tau} y(t) E[h(\theta)|\tau] dt dF_\tau(\tau) + v_0 \end{aligned}$$

Setting $n(\underline{\tau}) = 0$ is optimal. Now

$$\int_{\underline{\tau}}^{\overline{\tau}} \int_{\underline{\tau}}^{\tau} y(t)E[h(\theta)|\tau]dtdF_\tau(\tau) = \int_{\underline{\tau}}^{\overline{\tau}} \int_{t}^{\overline{\tau}} y(t)E[h(\theta)|\tau]dF_\tau(\tau)dt$$
$$= \int_{\underline{\tau}}^{\overline{\tau}} y(t)g(t)dF_\tau(t),$$

where $g(t) \equiv \frac{1}{f_\tau(t)} \int_t^{\overline{\tau}} E[h(\theta)|\tau]dF_\tau(\tau)$. This means

$$V_B(y, n) = \int_{\underline{\tau}}^{\overline{\tau}} y(\tau)r(\tau)dF_\tau(\tau) + \underline{\tau}y(\underline{\tau})E[c(\theta)] + v_0,$$

where $r(\tau) \equiv E[v(\theta)|\tau] - v_0 - \tau E[h(\theta)|\tau] + g(\tau)$.

Now the remainder of the proof follows the same steps as in the case where the regulator is judged only on consumer protection. □

Proof of Proposition 5. First consider the case where $h = 0$. For all $t < \tau^*$, there exists $\theta \in \Theta_t$ with $v(\theta) < v_0$, so $\min_{\theta \in \Theta_t} y(t)v(\theta) + (1 - y(t))v_0$ is maximized at $y(t) = 0$. Similarly, if $t > \tau^*$, then $v(\theta) \geq v_0$ for all $\theta \in \Theta_t$, so setting $y(t) = 1$ is optimal.

When the regulator is judged on consumer protection and firm costs, consider an approval rule y that gives some positive extent of approval to an effective type $\tau \leq \tau^*$. Then there exists $t \geq \tau$ such that $l(t) < v_0$, so $V_M(y) < y(t)l(t) + (1 - y(t))v_0 < v_0$. Since the policy which always rejects the firm achieves a payoff of v_0, such a y cannot be optimal.

Now fix an implementable approval rule y for which $y(\tau) = 0$ for all $\tau < \tau^*$. Note that there is a unique n which implements this y since $n(\underline{\tau}) = 0$ by IR and Assumption 2. For any θ such that $\tau(\theta) \geq \tau^*$, consider y' where

$$y' = \begin{cases} y, \text{ for } t < \tau(\theta) \\ 1 \text{ for } t \geq \tau(\theta) \end{cases}$$

and let n' be such that (y', n') is implementable. Now $y(t) = y'(t)$ and $n(t) = n'(t)$ are identical for $t < \tau(\theta)$. But by the envelope condition, $y(\tau(\theta))b(\theta) - c(\theta)n(\tau(\theta)) = y'(\tau(\theta))b(\theta) - c(\theta)n'(\tau(\theta))$. The preceding establishes that $J(y, \theta') \leq J(y', \theta')$ for all θ' such that $\tau(\theta') \leq \tau(\theta)$. Since preferences exhibit strictly increasing differences, note that for θ' such that $\tau(\theta') > \tau(\theta)$, $y(\tau(\theta))b(\theta) - c(\theta)n(\tau(\theta)) \leq y'(\tau(\theta))b(\theta) - c(\theta)n'(\tau(\theta))$. Therefore $J(y, \theta') \leq J('y, \theta')$ for every θ'. Since the choice of y and θ was arbitrary, this establishes that $J(y, \theta') \leq J(y^*, \theta')$ everywhere, i.e., y^* is the unique optimal undominated strategy. □

References

Banerjee, Abhijit, Sylvain Chassang, Sergio Montero, and Erik Snowberg, "A theory of experimenters," Technical Report, National Bureau of Economic Research 2017.

Bergemann, Dirk and Karl H Schlag, "Pricing without priors," *Journal of the European Economic Association*, 2008, *6* (2-3), 560–569.

Brown, Austin, Greg Rodriguez, Best Best, Krieger Tiffany Hoang, Hannah Safford, Gordon Anderson, and Mollie Cohen D'Agostino, "Federal, State, and Local Governance of Automated Vehicles," 2018.

Buckley, Ross P, Douglas W Arner, Robin Veidt, and Dirk A Zetzsche, "Building FinTech Ecosystems: Regulatory Sandboxes, Innovation Hubs and Beyond," *UNSW Law Research Paper*, 2019, (19-72).

Carpenter, Daniel, *Reputation and power: organizational image and pharmaceutical regulation at the FDA*, Vol. 137, Princeton University Press, 2014.

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

_ , "The political economy of FDA drug review: processing, politics, and lessons for policy," *Health Affairs*, 2004, *23* (1), 52–63.

Carroll, Gabriel, "Robustness and linear contracts," *American Economic Review*, 2015, *105* (2), 536–63.

Escobar, Juan, Qiaoxi Zhang et al., "Delegating Learning," Technical Report 2019.

Farber, Daniel A, "Uncertainty," *Geo. LJ*, 2010, *99*, 901.

Fenwick, Mark, Wulf A Kaal, and Erik PM Vermeulen, "Regulation tomorrow: what happens when technology is faster than the law," *Am. U. Bus. L. Rev.*, 2016, *6*, 561.

Frankel, Alex and Navin Kartik, "Muddled information," *Journal of Political Economy*, 2019, *127* (4), 000–000.

Frankel, Alexander, "Aligned delegation," *American Economic Review*, 2014, *104* (1), 66–83.

Friedman, Milton, "Frustrating Drug Review," *Newsweek*, Jan 1973, p. 49.

Gilad, Sharon, Moshe Maor, and Pazit Ben-Nun Bloom, "Organizational reputation, the content of public allegations, and regulatory communication," *Journal of Public Administration Research and Theory*, 2015, *25* (2), 451–478.

Guo, Yingni, "Dynamic delegation of experimentation," *American Economic Review*, 2016, *106* (8), 1969–2008.

_ **and Eran Shmaya**, "Robust Monopoly Regulation," *arXiv preprint arXiv:1910.04260*, 2019.

Henry, Emeric and Marco Ottaviani, "Research and the approval process: the organization of persuasion," *American Economic Review*, 2019, *109* (3), 911–55.

Heyward, James, Lily Olson, Joshua M Sharfstein, Elizabeth A Stuart, Peter Lurie, and G Caleb Alexander, "Evaluation of the Extended-Release/Long-Acting Opioid Prescribing Risk Evaluation and Mitigation Strategy Program by the US Food and Drug Administration: A Review," *JAMA internal medicine*, 2020, *180* (2), 301–309.

Janssen, Mathijs, "The Whole Truth? Generating and Suppressing Hard Evidence," 2019.

Kaplan, Sheila, "FDA faulted for failure to adequately monitor drug safety issues," Apr 2018.

McClellan, Andrew, "Experimentation and approval mechanisms," *Unpublished working paper*, 2017.

Miller, Henry I., "Opinion — Follow the FDA's Self-Interest," Oct 2018.

Moynihan, Donald P, "Extra-Network organizational reputation and blame avoidance in networks: The Hurricane Katrina example," *Governance*, 2012, *25* (4), 567–588.

Noll, Roger G, "Government regulatory behavior: A multidisciplinary survey and synthesis," 1982.

Olson, Mary K, "Regulatory agency discretion among competing industries: inside the FDA," *JL Econ. & Org.*, 1995, *11*, 379.

_ , "Firm characteristics and the speed of FDA approval," *Journal of economics & management strategy*, 1997, *6* (1), 377–401.

Richardson, Elizabeth, "Transparency in Clinical Research," Jan 2016.

Sunstein, Cass R, *Worst-case scenarios*, Harvard University Press, 2009.

Tetenov, Aleksey, "An economic theory of statistical testing," Technical Report, cemmap working paper 2016.

Thierer, Adam, *Permissionless innovation: The continuing case for comprehensive technological freedom*, Mercatus Center at George Mason University, 2016.

Tillio, Alfredo Di, Marco Ottaviani, and Peter Norman Sorensen, "Strategic sample selection," 2017.

Zhao, Ding, "Accelerated Evaluation of Automated Vehicles.," 2016.