

Approval Regulation and the Endogenous Provision of Confidence: Theory and an Analogy between Financial and Safety Regulation

Daniel Carpenter, Justin Grimmer and Eric Lomazoff*

October 26, 2009

Abstract

Recent years have witnessed renewed political and scholarly interest in consumer protection regulation; in the financial sector, legal scholars have proposed that safety regulation should govern financial markets (Warren 2007, 2008). In this paper, we ask whether the effects of safety regulation extend beyond the safety and quality of products to the initial confidence that consumers have about products. We model approval regulation, where a government regulator must approve the market entry of a product based upon observable, unbiased and non-anticipable experiments. We show that even if regulator and firm disagree only about quality standards, the disagreement induces the firm to provide more information about its product than it would in the absence of regulation. Put differently, purely first-order disagreements in regulation will generate second-order consequences (a less uncertain distribution of product quality). These second-order consequences are sufficient to generate first-order effects (more consumption of the product), even when users are risk-neutral. In other words, even if approval regulation produces little or no improvement in safety or quality, it still aggregates information useful to ‘downstream’ product users; these users will exhibit higher consumption and will more readily switch to superior products. In contrast with libertarian analyses of entry regulation, the model predicts that entry restrictions may be associated with greater product or service utilization (consumption) (Law 2003; Law and Marks 2009), as well as greater price sensitivity among consumers. Because contemporary cost-benefit analyses ignore these second-order effects, they are unlikely to capture the possible confidence effects of approval regulation.

*Earlier versions of this paper were presented at the Tobin Project Conference on Economic Regulation, White Oak, April 24-26, 2009; we thank Danny Goroff, Brigitte Madrian and David Moss for comments on previous versions. Carpenter [corresponding author] is Allie S. Freed Professor of Government [dcarpenter@gov.harvard.edu], and Grimmer and Lomazoff are Ph.D. Candidates, Department of Government, Faculty of Arts and Sciences, Harvard University. Comments, criticisms and corrections are most welcome.

Why not create a Financial Products Safety Commission charged with responsibility to establish guidelines for disclosure, collect and report data about the uses of different financial products, review new products for safety, and require modification of dangerous products before they can be marketed to the public? The agency could review mortgages, credit cards, car loans, and so on. It could also exercise jurisdiction over life insurance and annuity contracts. In effect, the FPSC would evaluate these products to eliminate the hidden tricks that make some of them far more dangerous than others and ensure that non pose unacceptable risks to consumers.

An FPSC would *promote* the benefits of free markets by ensuring that consumers can enter financial services markets confident that the products they purchase meet minimum safety standards.

– Elizabeth Warren, “Product Safety Regulation as a Model for Financial Services Regulation.”

A financial products safety commission could help fill in the gap, particularly in relationship to products being produced by and invested in by regulated entities. Each product would have to have a stated objective (e.g. in what ways was it helping manage and mitigate risk; what was the risk profile for whom the product was intended). Its risk characteristics would be identified, using conservative models which paid due attention to the failures previously noted. The Financial Products Safety Commission would evaluate whether products provided significant risk mitigation benefits of the kind purported by the product. There would be a presumption that there is no free lunch, i.e., that higher returns could only be obtained at the expense of greater risk; and a strong assumption against complex products, the full import of which are hard to analyze.

– Joseph Stiglitz, “The Financial Crisis of 2007/2008 and its Macroeconomic Consequences.”

1 Introduction

In the name of safety and quality, governments worldwide constrain the development, manufacture, marketing and utilization of products and services. The regimes of governance vary widely, from licensure of professional services to review of permit applications for construction or infrastructure renewal (such as dam or wetlands construction), and the extensive pre-approval requirements governing the introduction of drugs, medical devices and food additives. The ubiquity of these institutions is especially noteworthy. Elaborate regimes of professional licensure characterize not only modern-day economies worldwide, but also the historical past of economic development in North

America, Europe and Asia, among other regions. Mechanisms of pre-market approval for health products are observed not only in the United States, the European Union and Japan, but with increasing frequency (and rigor) in China, India, Brazil, Saudi Arabia and other less economically developed nations. In contrast to regimes of price regulation, moreover, institutions of quality and safety regulation have received much less analytic and scholarly attention.

Institutions of product and service regulation have recently attracted greater scholarly and policy attention, not least because scholars in law and economics (such as Warren and Stiglitz) have proposed that a form of ‘product safety’ regulation ought to be applied to financial markets. Among the essential ideas of these proposals are that a government regulator should prevent or limit the introduction of some products due to safety issues that their eventual consumers may not fully perceive. In the wake of the financial crisis of 2008 and 2009, these ideas and related notions and proposals have received considerable attention.¹

The idea of ‘product safety’ regulation for finance and other realms of economic industry is plausible and enticing, but also untried at the scope and scale for which its proponents are suggesting it. For this reason it deserves further investigation and scrutiny. How does product safety regulation work? What powers does a government regulator need to have in order to make safety and quality regulation work? What properties – what particular government powers and what particular sanctions and direct or indirect incentives – does quality and safety regulation entail? Relatedly, what properties might ‘product safety regulation for finance’ entail? Which properties of ‘consumer safety regulation’ are appropriable to other realms and which are not? Perhaps most important, does what is called ‘consumer safety regulation’ have implications for variables other than safety and product quality? In other words, is what we call ‘consumer safety regulation’ all about safety, or is in in part about something else?

¹In our reading, this discussion has been led by Warren, whose arguments are exceptional for two reasons: (1) she proposes a concrete analogy between a form of regulation that is widely known and familiar in one set of markets, and the absence of such institutions in financial and lending markets; and (2) she introduces the concept of ‘safety’ to discussions of financial regulation.

That ‘something else,’ we wish to claim, is confidence and experimental information. Although forms of safety and quality regulation can be thought to reduce the likelihood of product hazards or to mitigate against equilibrium fraud, these regulatory regimes do as much or more of their work by providing a certain kind of information about products. Both the quantity and quality of this information, we contend here, are far greater than would be provided in an unregulated marketplace. Accordingly, in this paper we draw upon existing formal models and historical studies of safety regulation to render one version of the safety-finance analogy more concrete. In this model the crucial powers of an approval regulator rest in two properties: (1) the ability to compel firms to experiment with new products before their introduction into the market, and (2) the ability to veto the introduction of new products, or to withdraw such products after their introduction. Property (2) is a form of regulation common to many forms of ‘entry regulation’; what makes this approval regulation regimes different is property (1), the robust connection between approval and the experimental generation of information used in the entry decision. This connection is observed in many licensure regimes (such that applicants must pass a series of courses or examinations, or perhaps a probationary period; Law and Kim 2005), in many permitting institutions (such that the proposed project must pass an environmental impact test or a cost-benefit criterion) and in pre-market approval institutions (such as pharmaceutical and medical device regulation worldwide).

More particularly, we present a mathematical connection between a stylized regulatory process and a stylized market. The regulation, the connection and the resulting market are all modeled, albeit with considerable simplification. In the regulatory process, a firm conducts experiments with its product to satisfy the higher quality/safety standards of a regulator, and the regulator sets an approval policy that skims off a subset of the products developed. We show that, even though the firm and the regulator disagree only about quality, their interaction produces information that reduces the uncertainty of both players. The results of this game are then aggregated into a single-mode, mixture distribution of product quality that is faced by consumers or users of the

product. In the resulting stylized market, a large number of ‘users’ consume the approved products from the distribution induced by approval regulation. In utilizing these products, consumers learn about the product’s quality. In the resulting user equilibrium, risk-neutral consumers are more likely to consume products (enter the market in the first place), and more likely to engage in optimal product switches once they do. First-order disagreements in regulation have second-order informational consequences, and these second-order consequences have first-order implications for consumers whose a priori utilities are unaffected by second-order considerations. If they exist, the second-order consequences of approval regulation regimes are unlikely to be captured by modern methods of policy analysis (Viscusi 1993; Boardman, et al., 1996).

Among the many limitations of our effort are three. First, our paper formalizes only several features of the analogy that Warren and other scholars draw between safety regulation and the possibilities for similar governance in financial and other markets. Mathematical models are notoriously poor at embedding the complex realities of public policy into analysis, and we make no claim here that all or even most of the salient features of outstanding proposals (including Warren’s, among many others) are incorporated in the present analysis.

Second, the model relies upon simplifying assumptions throughout, most notably the existence of a single firm at the regulatory stage and the exogeneity of consumers’ strategies to regulation.² While this separation of regulation from the utilization strategies of consumers may seem unduly constraining, it is instructive for several reasons. First, the connection drawn between regulation and consumer utilization is one that rests upon the communication of information produced in the regulatory process to ‘downstream’ users who might benefit from that information (such as doctors and patients who might benefit from the extensive knowledge produced in clinical trials conducted for FDA approval of a drug or medical device). This connection is accomplished in part through *labeling and other warning mechanisms*, which are little studied in political economy. Our

²This latter assumption is in part enabled by the large ‘extent of the market,’ hence any minor deviation by a single consumer from derived equilibria would have vanishing weight in the regulators’ calculations.

model makes it clear why these little-studied mechanisms of information transmission may in fact be so important. Second, our theoretical separation of regulation from consumption corresponds to the *de jure* separation of these phenomena in actual regulatory institutions. Many quality and safety regulators do not officially or explicitly consider downstream price and demand effects in their approval or licensure decisions (see for instance the provisions of the Federal Food, Drug and Cosmetic Act of 1938).

Third and finally, the model does not serve as a pure rationale for the creation of an ‘FDA’ or ‘phased-experiment’ system of trials for financial products or any other products. This is, not least, because we do not model all of the relevant costs of approval regulation regimes here. The really hard work of applied regulatory theory depends upon the careful translation of theoretical and normative considerations into policy proposals. In part, we intend this work to be undertaken by other scholars. In other ways, the model elaborated here may be informative for such efforts, as it points to particular features – the type of information generated and the incentives engendered by regulatory veto authority – that a policy proposal might entail.

On the positive side, we hope, the models here suggest that safety and quality regulation of the form that is observed in health and food products (and less strictly for licensure) may have broad benefits having less to do with safety and more to do with ex post product utilization. Our model offers results that cohere with recent analyses of product quality and safety regulation and their ex post effects in regulated markets (Law 2003, Law and Kim 2005, Law and Marks 2009). It also provides a formalization and a mechanism for the argument that Warren entertains in her essay, namely that, as Alexander Hamilton wrote two centuries ago in his essays on government and finance, the establishment of confidence should remain a critical focus of government policy.

2 A Model of Approval Regulation

Perhaps the essential characteristic of product safety regulation is in fact borrowed from political institutions: a *veto*, in this case a veto of the regulator over firms' market entry (or market continuation). We begin by modeling the centrality of this veto in the regulatory process, and in doing we offer a simple characterization of many features of licensing and product safety regulation in the modern world. What distinguishes our model from usual models of regulation is (1) the explicit modeling of entry vetoes, and (2) the centrality of experiment to the reduction of two-sided uncertainty.³

2.1 Game Structure

Informational Environment and Players: "No One Knows the Truth." There are two players: a firm ("the Firm") and a regulator ("the Regulator"). Both players are imperfectly informed about a parameter x of a product, which may be thought of as the product's "quality" or "safety." We assume that x follows a Beta distribution, or, formally ($x \sim \text{Beta}(\theta, n)$), where θ, n are positive integers, and $1 < \theta < n$. The first parameter of the distribution, θ , is the Firm's "type," that is, the estimated quality of its product.⁴ A higher type connotes a better expected product, while a lower type connotes a lower expectation of product quality. The Beta distribution of priors offers a natural interpretation of a set of n Bernoulli trials, of which θ resulted in success and $n - \theta$ in failure.⁵ The distribution is also flexible enough to accommodate a wide variety of "shapes" of the density function, as determined by θ and n . Given θ and n , the Beta distribution implies a prior mean θ/n and prior variance $\theta(n-\theta)/(n^2(n+1))$. The uncertainty over x can be resolved partially

³The elaboration follows Carpenter and Ting (2007), and many lateral proofs are presented there.

⁴ Since firms submit only one product in the model, we equate the type of the Firm with the type of its product.

It is possible to separate quality and safety and consider x to be their additive (hedonic) sum (e.g., "quality – safety," but a more compelling separation of the two variables this demands a much more complicated model with dual experimental strategies and possibly dual regulatory strategies. We leave this endeavor to another effort.

⁵The reputation games of Calvert (1987) and Alt, Calvert and Humes (1988) use a similar technology.

through observable experiments. Since the parameters affecting $E[x]$ change with experimentation, we distinguish between the Firm’s “initial” type θ and the numerical value of the first parameter of the Beta distribution by letting $m = \bar{\theta}$.

The product’s type is private information and known only to the Firm. We characterize product type as taking one of two possible values – “high” and “low” – hence $\theta \in \{\underline{\theta}, \bar{\theta}\}$ where $\underline{\theta} = \bar{\theta} - 1$. We will refer to $\underline{\theta}$ and $\bar{\theta}$ as the “low” and “high” types, respectively. The second parameter, n , is common knowledge. Let $p \in (0, 1)$ represent the Regulator’s prior belief of the high type ($\theta = \bar{\theta}$).

A crucial feature of our model, then, is that “nobody knows the truth” about the product’s quality in the sense that x is never known or estimated with full certainty. The Firm and Regulator possess only estimates of the truth, and information is asymmetric only to the extent that the Firm’s estimate (its “prior”) is better than the Regulator’s. The idea that even firms are somewhat uncertain as to the quality and safety of their products is a more accurate representation of regulation (and of other forms of social, political and economic reality). For this reason, it is, we believe, much more accurate and compelling than the simplistic representation of ‘incomplete information’ that has dominated the regulatory economics literature (Laffont and Tirole 1994).

Sequence of Play: Development First, then Regulation. The approval regulation game has as many as four periods ($t = 0, 1, 2, 3$), denoted by subscripts. These are divided into a *development* phase with up to three periods and possibly a *regulatory* phase of one period. The phases are distinguished by the mover: only the Firm moves in the development phase, and only the Regulator moves in the regulatory phase.

The approval regulation game begins in a “development phase.” In it, the Firm chooses an action $f_t \in \{S_F, W_F, E_F\}$ at $t = 1, 2$. S_F denotes a firm submission for approval, which ends the development phase and commences the regulatory phase the next period. W_F denotes the Firm’s withdrawal of a product from consideration, ending the game. Finally, E_F denotes a Firm-funded experiment to gather more data. An experiment is a single Bernoulli trial, which produces a publicly

observable result $e_t \in \{0, 1\}$ corresponding to failure or success, respectively. An experiment continues the development phase. The Firm cannot experiment past the second period and thus $f_3 \in \{S_F, W_F\}$. The final period thus represents something of a “fish or cut bait” choice for the firm. For convenience we let $e_0 = 0$, and use the term *experimental history* to refer to the set of experiments performed (up to two maximum).

At the beginning of the regulatory phase, the Regulator knows the Firm’s actions and experimental results, but does not know either x or θ . Based on this, she makes a review decision $r \in \{A, R\}$, where A and R denote acceptance and rejection of the Firm’s submission, respectively.

Revelation of Information: Calculation of Posterior Beliefs. Without experimentation, the model reduces to a simple signaling game with the Firm as the sender. With experimentation, both players update their expectations of x . The assumption of Beta-distributed priors makes the calculation of posterior beliefs very simple when both parameters of the distribution are known. For example, beginning with a prior of $\beta(\theta, n)$, two experiments producing $e_1 + e_2$ successes generate a posterior distribution of $\beta(\theta + e_1 + e_2, n + 2)$. Accordingly, $E[x \mid e_1, e_2] = \frac{\theta + e_1 + e_2}{n + 2}$ and $Var[x \mid e_1, e_2] = \frac{(\theta + e_1 + e_2)(n + 2 - \theta - e_1 - e_2)}{(n + 2)^2(n + 3)}$.

Utilities. The Firm receives x if the product is approved, and zero for rejection or withdrawal. Each experiment costs c_e , and a submission costs c_s , where:

$$c_e \in \left(0, \frac{m(m-1)}{(n+1)(n+2)(n+m-1)} \right] \quad (2.1)$$

$$c_s \in \left(0, \frac{m}{n+2} \right]. \quad (2.2)$$

These assumptions ensure that the low type is not prevented from experimentation or submission by exogenous costs alone (though it may choose not to do so in equilibrium). They also substantially simplify the analysis by eliminating some trivial equilibria.

The Regulator receives $x - k$ for an approved product, and zero otherwise. The parameter

k is therefore the divergence between the preferences of the Regulator and the Firm. It can be imagined to represent a “certainty equivalent” that a risk-averse or uncertainty-averse society, or a society fearing a product safety disaster, would demand from the firm in order for its product to be marketed. We assume that k satisfies:

$$k \in \left(\frac{m}{n}, \frac{m+1}{n+2} \right). \quad (2.3)$$

These bounds guarantee that some experimentation is necessary to generate a product satisfactory to the regulator, but that the two players can disagree over the desirability of marginal products.⁶

2.2 Equilibrium

We characterize Perfect Bayesian Equilibria (PBE) that satisfy a minor refinement. Let H_t represent the set of possible experimental histories prior to time t ; thus, $H_1 \equiv \emptyset$, $H_2 \equiv \{0, 1\}$, and $H_3 \equiv \{0, 1\} \times \{0, 1\}$. We use h_t to denote generic elements of H_t . The equilibrium has three elements.

1. The Firm’s strategy is the triple (ϕ_1, ϕ_2, ϕ_3) , where $\phi_t : \{\underline{\theta}, \bar{\theta}\} \times H_t \rightarrow \Delta(\{S_F, W_F, E_F\})$ for $t = 1, 2$, and $\phi_3 : \{\underline{\theta}, \bar{\theta}\} \times H_3 \rightarrow \Delta(\{S_F, W_F\})$ map types and experimental histories to probability distributions over submitting, withdrawing, and experimenting (where feasible).
2. The Regulator’s strategy ρ maps experimental histories, conditional upon a submission, into a probability of rejection. (Formally, $\rho : \bigcup_t H_t \rightarrow [0, 1]$).
3. The Regulator has beliefs μ mapping the experimental and submission history into a probability that $\theta = \bar{\theta}$. Formally, $\mu : \bigcup_t H_t \times \{\emptyset, S\} \rightarrow [0, 1]$. These beliefs must be consistent with Bayes’ Rule along the equilibrium path of play.⁷

⁶The upper bound on k ensures that a high type becomes acceptable to the Regulator after one experimental success, thus eliminating the rejection of “early” (*i.e.*, period 2) submissions as a dominant strategy. It also guarantees that the low type may become acceptable to the Regulator after two successful experiments, so that the Regulator’s problem is not merely one of “separating” the two types. The lower bound states that k is higher than both types’ *ex ante* expected product quality, so that neither will submit and be accepted in period 1.

⁷Along with h_t , this posterior distribution in turn induces a distribution over possible values of x .

Our analysis will utilize two other pieces of notation. First, we decompose $\phi_t(\theta, h_t)$ into probabilities of submitting, $\sigma_F(\theta, h_t)$, withdrawing, $\omega_F(\theta, h_t)$, and experimenting, $\eta_F(\theta, h_t)$. Because there are no more than three possible actions in each period for the Firm, $\sigma_F(\theta, h_t) + \omega_F(\theta, h_t) + \eta_F(\theta, h_t) = 1$ for all t .

Second, we denote the expected quality of a period t submission (given beliefs $\mu(\cdot)$) by:

$$x(\tilde{h}_t) = E[x \mid h_t, f_t = S] = \frac{\mu(h_t, S) + m - 1 + \sum_{i=1}^t e_{i-1}}{n + t - 1}. \quad (2.4)$$

This expression re-states the posterior mean of a Bernoulli process with Beta-distributed priors. The numerator gives the number of “successes” and the denominator gives the number of “trials,” and both incorporate the Regulator’s prior information about x and the experimental history.⁸

3 Main Results: Experimentation and Information Provision in the Approval Regulation Game.

The main model offers two types of equilibria— an *Early Submission Equilibrium* (ESE) and a *Late Submission Equilibrium* (LSE)— depending on how players react to a successful first-period experiment. This feature of the model introduces a layer of complexity, in that firm and regulator strategies are “history-dependent.”⁹

⁸As in Carpenter and Ting (2007), in order to complete the PBE, we must specify the Regulator’s response to out-of-equilibrium actions by the Firm. We assume that out of equilibrium, the Regulator’s beliefs about θ are given by $\mu(h_t, \cdot) = p$. Thus the Regulator essentially believes that each type is equally likely to deviate. For an out of equilibrium submission, (2.3) implies that $\bar{x}(h_t) < k$ for any $h_t \notin \{1, (1, 1)\}$. These beliefs are relatively innocuous. Any sufficiently “pessimistic” beliefs would support essentially the same equilibria as those identified here. If instead the Regulator’s beliefs were more “optimistic,” then for a larger set of histories the Regulator would accept all submissions. This would then induce all types to submit at those histories. It is easily shown that no equilibrium can be sustained in this manner.

⁹To simplify our presentation, we ignore “knife-edge” equilibria, which are not robust to small perturbations in parameter values. The remaining equilibria are the only ones in which the counter-intuitive strategy of withdrawing good products in the last period is not used. The proofs of Propositions 1 and 2 in Carpenter and Ting (2007) rule out a number of candidate equilibria. As the subsequent development shows, the predicted equilibria are unique for most, but not all, parameter values. This non-uniqueness does not affect the core results of the model that experimentation is generally increasing in k . We therefore eschew the imposition of another refinement.

3.1 General Features Shared by the ESE and LSE

The Regulator's decision problem occurs when the Firm submits. Clearly, the Regulator accepts a submission if its expected quality is greater than k , rejects if it is less, and is indifferent (and so may mix) otherwise. Thus:

$$\rho^*(h_t) = \begin{cases} 0 & \text{if } \bar{x}(h_t) > k \\ \in [0, 1] & \text{if } \bar{x}(h_t) = k \\ 1 & \text{if } \bar{x}(h_t) < k. \end{cases} \quad (3.1)$$

The Firm's choice will depend on its assessment of the value of experimentation. Let $v(\theta, h_t)$ denote type- θ 's continuation value from experimentation, conditional upon experimental history h_t . Clearly, $v(\theta, h_3) = 0 \forall \theta$. The Firm prefers submission over experimentation in period t if:

$$(1 - \rho^*(h_t)) \left[\frac{\theta + \sum_{i=1}^t e_{i-1}}{n + t - 1} \right] - c_s > v(\theta, h_t). \quad (3.2)$$

Finally, the Firm prefers submission or experimentation over withdrawal in period t if the expected payoff from either is non-negative.

3.2 Early Submission Equilibrium

An ESE is an equilibrium in which The Firm submits "early" (*i.e.*, in period 2) with positive probability. This can occur only if the Firm's first experiment is successful. There are two possibilities induced by the history $h_2 = 1$. These depend on the Firm's quality threshold, k , relative to \tilde{x} (defined in (2.4)), the expected period 2 quality conditional upon both types experimenting successfully in period 1.

If $k \leq \tilde{x}$, then the prior distribution of product technologies is favorable in the eyes of the Regulator. The Regulator is hence willing, at least in expectation, to accept the set of all successful first-period experimenters (including a product that would later be revealed to be of low type). Since additional experiments are costly and do not improve product quality in expectation, the

Firm submits whether its product is of high or low type, and all product submissions are accepted with certainty.

If $k > \tilde{x}$, then the Regulator would not accept the set of all first-period successes. However, she still wishes to accept the high type (by (2.3)), and can choose a rejection strategy to deter the low type from submitting with certainty. Early submission then requires that an initially successful high type prefers submission to continued experimentation. We label this the *Early Submission* (ES) condition:

$$c_s - c_e > \frac{m(m+1)}{(n+1)(n+2)}. \quad (3.3)$$

If the Firm's first experiment ends in failure, then by (3.2) she cannot submit. The continuation game then runs as follows. When c_s and c_e are sufficiently high, then it is possible that after $h_2 = 0$, the low type will prefer withdrawal to continued experimentation. This occurs under the following condition, which we label

Early Withdrawal (EW):

$$\frac{m-1}{n+1} \left(\frac{m}{n+2} - c_s \right) - c_e \leq 0. \quad (3.4)$$

When the Early Withdrawal condition holds, the Regulator correctly believes that only a high type would remain in the game after a failure, and thus has the same estimate of x as the Firm. Note that by (2.1) and (2.2), the truth of the Early Withdrawal condition implies the truth of the Early Submission condition.

Using these observations, the first result derives the ESE strategies.

(Early Submission Equilibrium) If $k \leq \tilde{x}$ or ES holds, then there exists an equilibrium given by:

For the Firm (type $\bar{\theta}$): $\eta^*(\bar{\theta}, \emptyset) = \eta^*(\bar{\theta}, 0) = 1$, $\sigma^*(\bar{\theta}, 1) = 1$, $\sigma^*(\bar{\theta}, (0, 1)) = 1$.

For the Firm (type $\underline{\theta}$): $\eta^*(\underline{\theta}, \emptyset) = 1$, $\sigma^*(\underline{\theta}, 1) = \begin{cases} 1 & \text{if } k \leq \tilde{x} \\ \frac{n(n+1)k - pm(m+1)}{(1-p)m(m-1)} & \text{if } k > \tilde{x}, \end{cases} \quad \eta^*(\underline{\theta}, 1) = 1 -$

$$\sigma^*(\underline{\theta}, 1), \eta^*(\underline{\theta}, 0) = \begin{cases} 0 & \text{if EW holds} \\ \frac{kn(n+1)(n+2)-p(n-m)m(m+1)}{(1-p)(n-m+1)(m-1)m} & \text{otherwise,} \end{cases} \quad \omega^*(\underline{\theta}, 0) = 1 - \eta^*(\underline{\theta}, 0), \omega^*(\underline{\theta}, (1, 0)) = 1, \sigma^*(\underline{\theta}, (0, 1)) = 1.$$

$$\text{For the Regulator: } \rho^*(1) = \begin{cases} 0 & \text{if } k \leq \tilde{x} \\ \frac{n-m+1}{n+2} + \frac{n+1}{m}c_e - \frac{n+1-m}{m}c_s & \text{if } k > \tilde{x}, \end{cases} \quad \rho^*(1, 0) = 1, \\ \rho^*(0, 1) = \begin{cases} 0 & \text{if EW holds} \\ 1 - \frac{n+2}{m} \left(\frac{n+1}{m-1}c_e + c_s \right) & \text{otherwise.} \end{cases} \quad \blacksquare$$

Proof Proofs of Proposition 1 and 2 appear in Carpenter and Ting (2007). All other proofs of Propositions are provided in the text, whereas proofs of Lemmata 1 and 2 appear in the Appendix.

■

For many parameter configurations and histories, the average quality of the Firm's submission is k , which makes the Regulator indifferent between rejection and acceptance. For histories $h_2 = 1$ and $h_3 = (0, 1)$, this allows R to choose rejection probabilities that make the low type indifferent between submitting and experimenting or withdrawing. In turn, the expected quality of submissions becomes exactly k . There are two cases in which the average quality exceeds k . The first, mentioned above, is the case where $h_2 = 1$ and $k < \tilde{x}$. The second occurs when the Early Withdrawal holds and $h_2 = 0$, which induces low types to withdraw and leaves only high types to generate the history $h_3 = (0, 1)$.

3.3 Late Submission Equilibrium

An LSE is an equilibrium in which there are no “early” submissions, that is, none at $t = 2$. This may occur either because the expected quality of successful experimenters at $t = 2$ does not warrant acceptance (*i.e.*, $k > \tilde{x}$), or because an initially successful experimenter would prefer to gather more information (*i.e.*, the Early Submission condition is violated). The existence of the LSE is assured when the ESE does not exist, and furthermore does not depend on the values of c_s or c_e . The LSE requires only that $k > \tilde{x}$, while the ESE requires that either $k \leq \tilde{x}$ or the Early Submission condition hold. Thus when $k \leq \tilde{x}$, the model uniquely predicts the ESE. And when $k > \tilde{x}$ and the

Early Submission condition does not hold, the model uniquely predicts the LSE. If $k > \tilde{x}$ and the Early Submission condition holds, then both equilibria exist.

Analogously to Proposition 3.2, the next result derives the LSE strategies.

(Late Submission Equilibrium) If $k > \tilde{x}$, then there exists a equilibrium given by (3.1)-(3.2)

and:

For the Firm (type $\bar{\theta}$): $\eta^*(\bar{\theta}, \emptyset) = \eta^*(\bar{\theta}, 0) = \eta^*(\bar{\theta}, 1) = 1$, $\sigma^*(\bar{\theta}, (0, 1)) = \sigma^*(\bar{\theta}, (1, 0)) = 1$.

For the Firm (type $\underline{\theta}$): $\eta^*(\underline{\theta}, \emptyset) = \eta^*(\underline{\theta}, 1) = 1$, $\eta^*(\underline{\theta}, 0) = \begin{cases} 0 & \text{if EW holds} \\ \frac{kn(n+1)(n+2)-p(n-m)m(m+1)}{(1-p)(n-m+1)(m-1)m} & \text{otherwise,} \end{cases}$
 $\omega^*(\underline{\theta}, 0) = 1 - \eta^*(\underline{\theta}, 0)$, $\sigma^*(\underline{\theta}, (0, 1)) = 1$, $\sigma^*(\underline{\theta}, (1, 0)) = \frac{kn(n+1)(n+2)-pm(n-m)(m+1)}{(1-p)(m-1)(n-m+1)m}$.

For the Regulator: $\rho^*(1) = 1$, $\rho^*(1, 0) = 1 - \frac{n+2}{m}c_s$, $\rho^*(0, 1) = \begin{cases} 0 & \text{if EW holds} \\ 1 - \frac{n+2}{m} \left(\frac{n+1}{m-1}c_e + c_s \right) & \text{otherwise.} \end{cases}$

■

In both the ESE and LSE, the average quality of submitted products is usually k , which makes the Regulator indifferent between rejection and acceptance. The only exception occurs when $h_3 = (0, 1)$ and the Early Withdrawal condition holds, since low types will have withdrawn from experimentation. This results in submission quality higher than k and acceptance by the Regulator.

As an example of an LSE game history, suppose that the first experiment of a high type firm succeeds. The Regulator does not accept submissions at $t = 2$, so the Firm experiments again and submits regardless of the outcome. If $e_2 = 1$, then the Regulator accepts. If $e_2 = 0$, then the Firm's submission is still acceptable to the Regulator, but the low type's submission has an expected quality below k . Thus, the Regulator mixes between acceptance and rejection.

3.4 Implications of Approval Regulation Equilibria

A critical result from both Early Submission and Late Submission Equilibria is that they result in the production of more information about the product than would occur if the Firm could enter the marketplace without a regulatory veto. This is true despite the fact that both players are risk-

neutral and indeed uncertainty-neutral; neither Firm nor Regulator care directly about information (the option value of experimentation is not informational but instead a “gambling” option value, that is, contingent upon the possibility of superior results). The intuition here is that the Firm can only enter the market with the regulator’s assent, and there is no other way to gain the regulator’s assent but to prove the product’s case. Since the Firm’s threshold for a desirable product is lower than the Regulator’s (by the scalar term k), the Firm conducts more experiments with its product than it would if it were doing R&D on its own. This result is summarized and demonstrated in Lemma 1.

Lemma 1: Second-Order Stochastic Dominance of Regulated-Induced Product Distribution

Denote the unregulated product distribution by $E_U[x]$ and the distribution produced in the ESE or LSE by $E_R[x]$. For all $k > 0$, $E_R[x]$ has second-order stochastic dominance over $E_U[x]$.

Proof appears in the Appendix.

Intuitive Summary. the approval regulation model as such offers several useful lessons for mapping into the regulated industry domain. First, costly activity by The Firm can serve multiple purposes, thus distinguishing our model from standard costly signaling models. Here experimentation and submission are both signals of type, but experimentation also endogenously generates additional information. In other words, information has two dimensions of value (private information and common information). This information, in addition to type, determines submission strategies. Second, because acceptance with certainty would induce all types to submit, the expected quality of submitted and accepted products will often be exactly the Regulator’s reservation value of k . The equilibrium acceptance rate induces the low type to submit at a point of induced indifference between submitting and experimenting (or, in the final period, between submitting and withdrawing). The Regulator thereby benefits from its gatekeeping power, as $k > \frac{m}{n}$. Thus the approval regulation process allows the Regulator to “skim” the best products from a population

that is *ex ante* unacceptable.¹⁰

4 Cumulation of Regulatory Decisions – An Ex Ante Product Distribution for Consumers

We now construct a thought experiment by imagining that the approval regulation game in the previous section is repeated and that the induced product distribution is the sum of many repetitions of the game. In so doing we construct a mathematical map between the regulatory process and a market, but do not model both jointly.¹¹

Lemma 2: Cumulation of the Regulation-Induced Product Distribution into a Continuous Mixture Distribution

Let $G(\cdot)$ represent a non-degenerate probability distribution with primitive g . Take the product distributions $G_U \equiv G(E_U[x])$ and $G_R \equiv G(E_R[x])$, and express π_i as a convex mixing distribution on $[0,1]$. Consider a countable set of such distributions, each i.i.d., with $\zeta \in \mathbb{N}^+$ an index variable, and with each mixture parameter α expressed as the integer ratio $\frac{a'_\zeta}{b'_\zeta}$ where $a'_\zeta \leq b'_\zeta \forall \zeta$. Assume the moments of α follow those of a Beta distribution. For N_ζ repetitions of the product distribution, construct the integral distributions $G_U^\Sigma = \liminf_{\zeta \rightarrow \infty} \sum_\zeta \left\{ \pi_\zeta^U(G^U)/N_\zeta \right\}$ and $G_R^\Sigma = \liminf_{\zeta \rightarrow \infty} \sum_\zeta \left\{ \pi_\zeta^R(G^R)/N_\zeta \right\}$. Assume further that the integral distributions are square integrable, i.e., for both regulated and unregulated distributions, $E \left[\liminf_{\zeta \rightarrow \infty} \sum_\zeta \left\{ \pi_\zeta(G)/N_\zeta \right\} \right]^2 < \infty, \forall \zeta > 0$. Then G_R^Σ and G_U^Σ are continuous-parameter distributions having support on \mathfrak{R} . Each distribution is approximately Normal, and G_R^Σ SOSD G_U^Σ .

Lemma 2 is somewhat technical but essentially says that if the approval regulation game were averaged over many independent plays, the resulting distribution of product quality would be

¹⁰The flip side of this mixed strategy is that, because of asymmetric information, the Regulator commits both Type I and Type II errors. Errors ensue because submissions are pooled and the Regulator can do no better than to mix between acceptance and rejection. This result contrasts with a decision-theoretic world in which the Regulator (or some other representative agent) knows θ and experiments on her own. In that case no errors would occur.

¹¹This remains as a research agenda, but see the Conclusion for a discussion of very real barriers of mathematical tractability for doing so, and see the Introduction for the difficult but pointed research questions that arise from thinking about this real-world map as a next modeling and empirical step.

continuous and approximately Normal. In addition, the stochastic dominance relations in Lemma 1 would be preserved under this averaging. Invocation of Lemma 2 permits the result of the approval regulation game to be applied to learning and utilization problems that have continuous representations of uncertainty as well as discrete representations.

5 A Generalized Dynamic Model of Utilization under Uncertainty in an Approval-Regulated Market

We now turn to model utilization or consumption in the “market” for the products approved by regulation. In so doing, we present a very stylized version of utilization or “consumption” in which a large number of human agents each uses one product at a time, switching among products to find the best one available.¹² The central fact governing the agents’ consumption decisions is uncertainty – the different products can be rank-ordered by the prior estimates of their quality (their “initial appearances”) but these are only best guesses subject to error. Beyond the initial best guesses, human agents’ uncertainty about products can be reduced only through experience, such that the regulated product is an experience good.

The uncertainty is expressed as a dynamic stochastic process, and here we use the canonical version, a Wiener process or “Brownian motion.”¹³ One can think of a Brownian motion as an all-purpose random process whose independent movements in continuous time occupy a continuous state space. For more than a century (since Louis Bachelier’s *Théorie de la spéculation* (1900)), Brownian motion has been used to model the movement of asset prices, and a general example of applications to a wide variety of investment decisions occurs in Dixit and Pindyck (1994). Brownian motion is useful for our purposes (and many others) because it offers a simple, additive and linear

¹²We eventually assume that there are many human agents, such that the market is “large,” or at least sufficiently large that coordination among agents becomes prohibitively difficult. This lack of ability to coordinate – say, for consumers to conduct their own tests of products – can be thought of as contributing to the need for a regulator.

¹³For a similar problem with more complex ‘spectrally negative’ Levy processes that include discontinuous jumps, see Carpenter and Grimmer (2009). We aim eventually to embed a multi-armed bandit model in this Levy process framework.

expression of uncertainty.

We begin with a single human agent (hereafter “Agent”), who can be thought of as an end-user or consumer of the products produced by the Firm and governed by the Regulator. The Agent observes the unfolding of realized product value on a space Λ (with elements or experimental realizations λ), which is structured by a set of σ -algebras \mathfrak{F} , and a probability measure φ . In addition, \mathfrak{F} can be ordered and expressed as a filtration $(\mathfrak{F}_t)_{0 \leq t \leq \infty}$, which is a family of σ -algebras that is increasing in its index, hence $\mathfrak{F}_s \subset \mathfrak{F}_t$ if $s \leq t$. The filtration sequentially collects and orders all realizations $\lambda = \lambda_t$ on a time dimension from 0 to t . The collection $(\Lambda, \mathfrak{F}, \mathfrak{F}_t, \varphi)$ constitutes a filtered probability space. This filtered probability space supports a standard one-dimensional Brownian motion $Z(t)$, and we assume that a set of “usual hypotheses” hold. These hypotheses are standard in the analysis of stochastic differential equations (see Protter (2005: Chapter I, esp. pp. 34-36) for a clear explanation).

Let the incumbent product under consideration be indexed by i (where the countable set of available products is indexed by j), and let the consideration time for product i be given by t_i . A fallback product with known value is denoted by $i = 0$, whereas all incumbents and uncertain products (“arms” of the bandit) are denoted by $i \geq 1$. We suppose that each product is characterized by a *quality* parameter. In a health example, this will reflect the product’s effectiveness in treating the disease; in a financial example, this will reflect its contribution to the Agent’s (the investor’s or asset holder’s) welfare. A product’s quality is a draw from a normal distribution, $\beta_i \sim \Phi(b, v_b)$ where $\Phi(b, v_b)$ represents the normal distribution with mean b and variance v_b . The actual value of β_i is unknown to the Agent, but is learned from the experience of utilization; only one product can be utilized at a time.¹⁴

¹⁴As long as convex bundles of products impart some randomness relative to previously utilized components, the results of this one-product-at-a-time model will generalize easily.

5.1 Continuous Time Evidence of Quality

The Agent collects continuous-time evidence about a product’s quality according to Brownian motion with drift, where the drift is determined by the (unobserved) quality (β_i) of the case. Formally, the Agent observes and experiences the realized value of the product X_i , which evolves according to the following stochastic differential equation.

$$dx_j(t) = \beta(x_j(t)) dT_j(t) + \xi(x_j(t)) dz_j(T_j(t)); t > 0$$

where T_j is the learning or utilization time for the j th product, z_j is a standard normal distribution with mean zero and variance t_j . As we show below, the parameter ξ encodes the amount of information Equation 5.1 contains for the regulator: if $\xi_i = 0$ then the regulator can immediately infer the quality of the case by examining the slope of Equation 5.1 and as $\xi \rightarrow \infty$ the SDE contains no information about a product’s quality.

5.2 Estimating Quality from Evidence

Given that the Agent only observes $X_i(t)$ we first prove that the learning problem is identified: the Agent is able to disentangle the contribution of the quality of the case to $X_i(t)$.

Identification of Learning Problem and Sufficient Statistics. We assume fixed coefficients and adopt the technology of Herman Chernoff (1968), who presents closed-form Bayes posteriors of a Brownian motion with drift.¹⁵ Without loss of generality, then, for any $X_i(t_i)$, the history of $X_i(t)$, $\mathbf{H}_i(t_i)$ can be expressed by its sufficient statistics, namely the dual $(t_i, X_i(t)^*)$. Let $\Pi \in (U, R)$ denote whether the product in question is Regulated or Unregulated. Then,

$$\text{Posterior Mean} \equiv E_{xt}(\beta_i) = \hat{\beta}_{it} = \frac{b/v_b^\Pi + x_i/\xi^2}{1/v_b^\Pi + t_i/\xi^2} \quad (5.1)$$

¹⁵By scale invariance of the Brownian diffusion (Karatzas and Shreve 1991: 66-71), the usual operators and Lemmata of Ito calculus can be applied straightforwardly to these posterior quantities.

$$\text{Posterior Variance}(\hat{\beta}_i) \equiv V_{\hat{\beta}_i}^{\Pi}(t_i) = \frac{1}{1/v_b^{\Pi} + t_i/\xi^2} = \frac{1}{[v_b^{\Pi} (\theta^{\Pi}, n^R)]^{-1} + \xi^{-2}t_i} \quad (5.2)$$

5.3 Filtered Evidence and Value Functions

The Agent seeks to define an optimal continuation rule for the *filtered evidence process* found by combining Equations 5.1 and 5.2, $\hat{\beta}_i(t_i)$. The Agent faces a convex function $\hat{\beta}_i(t) \times t \mapsto \Psi(\hat{\beta}(t)_i, t_i)$, that is twice differentiable with respect to both $\hat{\beta}(t)_i$ and t . This function is a map from the current state of the filtered evidence process and time to the value experienced by the agent. For any particular product, the Agent wishes to maximize

$$E \int_0^{\infty} e^{-\delta t_i} \bar{h}(i(t_i), x_{i(t_i)}, q_i) dt_i$$

where δ is a discount factor interior to the unit interval and $\bar{h}(x_i)$ is a running payoff function or flow value, realized when observed quality is x_i and q_i is the per-unit-time (dt) cost of utilization, with $\frac{dq_i}{dt} \equiv 0$, $\frac{d\hat{\beta}_i(t_i)}{dt} \equiv 0$ and $\frac{\partial \bar{h}(x_i)}{\partial q_i} < 0$. For the moment (until consideration of Proposition 6 below), assume that utilization cost is identical across products $q_i \equiv q_j \equiv q, \forall i, j$.¹⁶ In the health example, \bar{h} can be considered as the patient's realized and experienced health; in many investment and financial examples this is construed as a dividend. For the following analysis we will replace x_i with $\hat{\beta}_i$, without loss of generality due to the scale-invariance property of $X_i(t_i)$.

5.4 Utilization/Consumption Objective

The Agent can utilize only one option at a time (the incumbent) but faces a countable sequence of alternatives, including one that has certain value (a “fallback” option). For each incumbent, the Agent's objective is to define an optimal rule to stop $\hat{\beta}_i(t_i)$ in order to maximize her expected reward, which is given by

¹⁶The utilization cost is of course a commodity price, and we are assuming its exogeneity to user behavior and regulation here. Endogeneity to user behavior would necessitate a dynamic stochastic equilibrium as in Jovanovic (1979), who does not model regulation or an alternative battery of products/jobs.

$$J(x, M; \mathfrak{S}, \tau) = E_x \left[\int_0^\infty e^{-\delta t} \bar{h} \left(i(t), \hat{\beta}_{i(t_i^*)}^*(t_i^*) \right) dt_i + M_i e^{-\delta \tau} \right]$$

where M is an idiosyncratic termination payoff which is static, positive and known with certainty, and τ is a stopping time at which point the Agent hypothetically switches to the known fallback option (described below) or a product with its certainty equivalent. The quantities β_i^* and t_i^* are the case's quality and current time under the optimal learning policy.

The solution to a problem of this nature is well known and entails the human agent's calculation of a dynamic allocation index (DAI) or "Gittins index." Intuitively, the Gittins index of an uncertain product is the minimum certain reward that an agent would choose over that product, given everything that the agent knows about the uncertain product. For any given β_0 , let $M = E \int_0^\tau e^{-\delta t_i} \beta_0 dt_i$ stand for the expected terminal reward associated with choosing the fallback option ($i = 0$) until τ .¹⁷ For the incumbent (i th) arm, and a (possibly optimal, possibly biased) estimate of quality $\hat{\beta}_t(t_i)$, the Agent's optimal strategy is to establish a "follow the leader" policy associated with Gittins index Q_i , which establishes the incumbent product according to the following monotonic optimization schedule,

$$Q_i \left(\hat{\beta}_t(t_i) \right) = \inf \{ \kappa \in \mathfrak{R}^+ \mid \sup_{\tau > 0} \frac{E \int_0^\tau e^{-\delta t_i} \bar{h} \left(i(t_i), \hat{\beta}_{i(t_i)} \right) dt_i}{E \int_0^\tau e^{-\delta t_i} \beta_0 dt_i} = \kappa \}$$

where the supremum is over all \mathfrak{S}_i -stopping times that are positive a.s., and at least one is not necessarily finite.

Given a calculable Gittins index for any option, then the agent's optimal strategy is to choose the option with the maximum Gittins index (Karatzas 1987; Mandelbaum 1987). This, then, is how a rational agent would dynamically utilize a countable battery of products, where one has known value. Again, this technology is well-known and is a special case of bandit learning developed in

¹⁷It is assumed, without loss of generality to the model, that the Agent starts utilizing a product with unknown value.

statistical decision theory (Gittins 1979; Banks and Sundaram 1992).

5.5 The Optimal Stopping Rule and Its Properties

Because each product is associated with a prior $\beta_j^0 = \hat{\beta}_t(t_j = 0)$, the battery of products can be arranged ordinally (and monotonically with respect to their prior value). As has been amply demonstrated in mathematical statistics (Karatzas 1987; Mandelbaum 1989), the Agent’s optimal policy will be to utilize the product with highest initial estimated quality and pursue an optimal switching policy from there. Ties have measure zero and the agent can be assumed to randomize among tied products. The only additional assumption required to pin down the model is the existence of a termination penalty $\chi_j > 0$ for any abandoned product, such that once the j th option is abandoned, its prior for any possible future consumption will be reduced by χ_j .

The the Agent’s sequential problem is equivalent to following the leader according to an optimal stopping policy. Begin utilizing the product with the highest Q_i , then abandon this incumbent if and only if, and when and only when, an optimal stopping criterion is satisfied.

Using the scale-invariance of diffusions (Karatzas and Shreve 1991: 66-71), the Agent’s optimal policy will be to observe the first passage of the evidence process $\hat{\beta}(t)_i$ through a border that encodes the tradeoff between continuation of experience with the incumbent product and the value of switching (with partial irreversibility since $\chi_j > 0$) to the next best product, where “next best” simply encodes the best prior among the non-incumbent options: $\sup_{1 \leq j \leq d} \beta_j^0$. Then among the countable set of alternative products, the agent can consult and compute a Hamilton-Jacobi-Bellman equation for each product j :

$$\delta\Psi(x) = \max_{\{1 \leq j \leq d\}} \frac{\partial\Psi(x)}{\partial\hat{\beta}_j} + \frac{\partial\Psi(x)}{\partial t_j} + V_{\beta_j}^{\Pi}(t_j)^2(x_j(t_j)) \frac{\partial^2\Psi(x)}{\partial\hat{\beta}_j^2} + \bar{h}(j, x_j) + o(t_j)$$

where $o(t_j)$ denotes “vanishing” terms of order greater than t_j , that is, terms that converge to zero faster than t_j does. A critical result in stochastic process theory, the Shiryaev separation

theorem (Shiryaev 1978: 100), permits the Bellman equation for the optimal stopping problem to be expressed in terms of an infinitesimal generator without the flow payoff term \bar{h} . After applying Ito's Lemma, independence, and the pure martingale property, dividing through by the differential dt_i and taking limits as the differential vanishes, the infinitesimal generator \mathcal{L} for the incumbent evidence process $\hat{\beta}_i(t_i)$ can then be expressed as:

$$(\mathcal{L}_i^x \Psi)(x_i) = \Psi_{\hat{\beta}_i}(x_i, t_i) + \Psi_t(x_i, t_i) + \frac{1}{2} V_{\hat{\beta}_i}^\Pi(t_i)^2 \Psi_{\hat{\beta}_i \hat{\beta}_i}(x_i, t_i)$$

Evaluating $\mathcal{L}_i^x \Psi(x_i)$ according to the Shiryaev conditions (smooth pasting and value matching; Shiryaev 1978) results in elimination of the Ψ_t term and a uniquely optimal first-passage time policy. We have proved:

Proposition 3: Optimal Stopping Barrier for Each Incumbent Product

The Agent switches products when and only when, and if and only if, $\hat{\beta}(t)_i$ passes for the first time through the optimal stopping barrier established by the availability of the next-best product,

$$\gamma^*(t) = \delta Q_j + \frac{1}{2\xi^2} \Psi_{\hat{\beta}, \hat{\beta}}(\hat{\beta}(t)_i, t) V_{\hat{\beta}_i}^\Pi(t_i)^2 \tag{5.3}$$

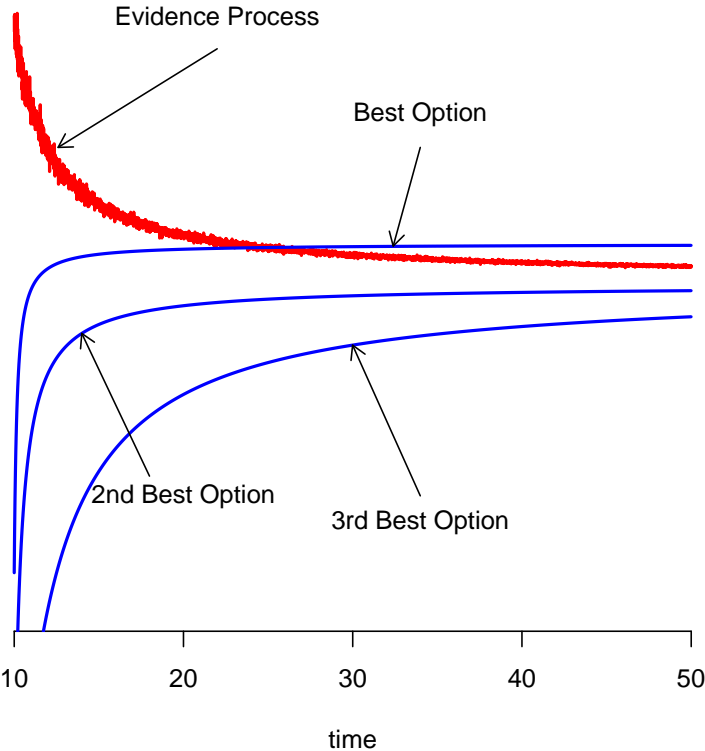
where $\Psi_{\hat{\beta}\hat{\beta}}(\hat{\beta}(t)_i, t)$ is the second partial derivative of the value function Ψ with respect to the filtered state variable $\hat{\beta}$, given a realization of $\hat{\beta}$ at time t . For the terminal product with known value, replace Q_j with M .

This border represents the optimal tradeoff between continuing with the incumbent product (and delaying the utilization of the next-best product) and switching instantly. If the Agent delays utilization of a case, she receives more information, reducing the value of $V_{\hat{\beta}_i}^\Pi(t_i)$, which under the optimal stopping policy declines quadratically with the length of the Agent's experience.

Figure 1 visualizes the filtered evidence process and the optimal stopping barrier from Propo-

sition 3. The horizontal axis represents time, the vertical axis is the utility provided to the Agent, the red-line is the filtered evidence process for one product, and the blue/purple lines represent families of optimal stopping barriers for the set of alternative products. The border slopes upward, as the value of more information decreases over the course of the regulatory history. A case is approved only if its evidence crosses the boundary, which occurs at the right-hand side of Figure 1.

Figure 1: Evidence Process and Barriers from Proposition 3



We are now in a position to advance our principal theoretical result. Before doing so, we simply restate two central features of our elaboration so far – that the payoff structures of the approval regulation game and the market are linear, additive and entail only the “first moment” or expected value of the random variables involved. In other words, not one actor in the models we have

presented – neither firm, nor Regulator, nor Agent (end-user) – is “risk-averse,” and none of these actors directly values higher moments or orders of the uncertainty (such as the variance or skew) they experience.¹⁸

Proposition 4: Higher Switching Rate to Superior Products by Agents under a Regulation-Induced Quality Distribution when the Unregulated Product Has Superior Priors

Consider any two alternative products $j = 1, 2$ where the first arrives to the market via approval regulation (is Regulated) and the second arrives without approval regulation (is Unregulated). Assume

(1) that the products are of identical value and that both are superior to the incumbent product, such that $[\beta_{j=1,R}^{\hat{\beta}_R(0)} \equiv \beta_{j=2,U}^{\hat{\beta}_U(0)} > \beta_i^{\hat{\beta}_i(0)}]$,

(2) that the alternative products generate identical coefficient histories,

$$H \left(X_{1,R} (t_j) | \beta_{1,R}^{\hat{\beta}_R(0)} \right) \equiv H \left(X_{2,U} (t_j) | \beta_{2,U}^{\hat{\beta}_U(0)} \right), \forall t_j$$

(3) and that the initial quality estimate of the unregulated alternative product is higher than the initial quality estimate of the regulated alternative ($\hat{\beta}_U(0) > \hat{\beta}_R(0)$).

Then the agent still switches more quickly to the regulated product (R) at $t_j = t_{\text{switch}}$ unless the prior difference is sufficiently large that

$$\hat{\beta}_U(0) - \hat{\beta}_R(0) > \frac{\xi^2 + v_b t_{\text{switch}}}{\xi^2} \left[V_{1,R}^{\Pi} (t_{\text{switch}})^2 - V_{2,U}^{\Pi} (t_{\text{switch}})^2 \right]$$

Proof: Note first that by the Borel-Cantelli Lemma, the agent will eventually switch to both

¹⁸Of course, the Agent in the diffusion bandit model does embed second-order considerations in dynamic utilization, but this is induced because χ_j is strictly positive.

products, as $\liminf_{t \rightarrow \infty} \Pr \left[\hat{\beta}_i(t) < \gamma^*(t)[\beta_U] | \beta_U > \beta_i \right] = \liminf_{t \rightarrow \infty} \Pr \left[\hat{\beta}_i(t) < \gamma^*(t)[\beta_U] | \beta_U > \beta_i \right] = 1$.

There are, in other words, no possible Type II errors in this model. The question then becomes whether the agent switches more quickly to the regulated or unregulated product. Because the prior quality estimates differ, the posterior quality estimates converge only in the asymptote:

$$\hat{\beta}_t^{\hat{\beta}_U(0)} \left\{ H \left(X_{1,R}(t) | \beta_{1,R}^{\hat{\beta}_U(0)} \right) \right\} > \hat{\beta}_t^{\hat{\beta}_R(0)} \left[H \left(X_{2,U}(t) | \beta_{2,U}^{\hat{\beta}_R(0)} \right) \right], \forall t_j < \infty$$

For any single barrier, the difference between these estimates can be described as

$$\begin{aligned} \hat{\beta}_t^U [H] - \hat{\beta}_t^R [H] &= \frac{\left(b + (\hat{\beta}_U(0) - \hat{\beta}_R(0)) \right) / v_b + x_{1t} / \xi^2}{1/v_b + t/\xi^2} - \frac{b/v_b + x_{2t} / \xi^2}{1/v_b + t/\xi^2} \\ &= (\hat{\beta}_U(0) - \hat{\beta}_R(0)) \xi_i^2 (\xi^2 + v_b t_j)^{-1} \end{aligned}$$

Let Q be the termination payoff for the incumbent, which is simply the Gittins index value for the next-best product (or M if only the fallback option remains). Then define a stopping time of the regulated product by

$$t_{\text{stop, R}} = \inf \left[t, \quad \text{s.t.} \quad \hat{\beta}_{1,R,t} < \gamma(t) \mid \beta_{1,R} < \delta Q \right]$$

and similarly define a stopping time of the unregulated product by

$$t_{\text{stop, U}} = \inf \left[t, \quad \text{s.t.} \quad \hat{\beta}_{2,U,t} < \gamma(t) \mid \beta_{2,U} < \delta Q \right]$$

By assumptions (1) - (3) of the Proposition, $t_{\text{stop, U}} < t_{\text{stop, R}}$ if the barriers are identical. Yet by Lemma 1, the prior variance of the regulated product will be lower due to regulation-induced confidence. Hence a sufficient condition for the regulated product to gain quicker approval can be derived as

$$\left[\delta Q - V^{\Pi,R}(t_{\text{stop}, R})^2 \right] - \left[\delta Q - V^{\Pi,U}(t_{\text{stop}, R})^2 \right] \geq \hat{\beta}_{t_{\text{stop}, R}}^U[H] - \hat{\beta}_{t_{\text{stop}, R}}^R[H]$$

Note that the comparison is conducted at $t_{\text{stop}, R}$ for both products. Set $t_{\text{switch}} \equiv t_{\text{stop}, R}$ and solve to get

$$\hat{\beta}_U(0) - \hat{\beta}_R(0) \leq \frac{\xi^2 + v_b t_{\text{switch}}}{\xi^2} \left[V^{\Pi,U}(t_{\text{switch}})^2 - V^{\Pi,R}(t_{\text{switch}})^2 \right]$$

which was to be shown.

This result is actually quite general because it relies upon the equivalence of stochastic histories – the Agent’s accumulated data on the regulated and unregulated products – and because the stopping times used in the proof are arbitrarily chosen (and, given the equivalence of histories, substitutable across any two products). More complicated (but no less rigorous) proofs are available using stochastic integration.¹⁹ Moreover, the equivalence of stochastic histories can be relaxed slightly and analysis of the model yields a similar result (see Lemma 3 in the Appendix). We are now in a position to state a more general result, which follows as a corollary from Proposition 4.

Proposition 5: Higher Switching Rate to Superior Products by Agents under a Regulation-Induced Quality Distribution

Consider any two products $j = 1, 2$ satisfying assumptions (1) and (2) of Proposition 4. The agent operating under the optimal dynamic allocation policy will always switch more quickly and with higher probability to superior products under regulation than under an unregulated product distribution.

Proof: Immediate. With $\hat{\beta}_U(0) - \hat{\beta}_R(0) = 0$, then by Lemma 1 the condition of Proposition 4 is satisfied almost surely.

Propositions 4 and 5 offer one other lesson that has some mathematical interest but is also

¹⁹We say no less rigorous because the equivalence of stochastic histories implies a control for the measure-zero event “all other things held constant.”

of useful relevance to more practical policy considerations and is, furthermore, testable, perhaps both experimentally and observationally. Even small reductions in the prior uncertainty of the set of potentially utilized products have great effects because the reduction of uncertainty behaves according to quadratic convergence.²⁰ So even a ‘small amount’ of induced confidence from regulatory institutions can have a great effect upon utilization of the product, because initial data has the greatest marginal effect.

Implications for Consumption and Price Sensitivity of Agents. The results in Propositions 4 and 5 have an interesting corollary for the role of price in a regulated market. The enhanced pre-market information about a product may induce a behavior where consumers are more price-sensitive because they are less sensitive to *a priori* uncertainty about product safety or quality. A full investigation of this dynamic would require the technical construction of the price-elasticity of demand in the bandit model; we instead proceed with the marginal change in the dynamic agent’s willingness to pay for the incumbent product for another period (dt) for each change in price.

Proposition 6: Regulation and Price Sensitivity of Switching Rate to Superior Products

Consider again any two products $j = 1, 2$ satisfying assumption (2) of Proposition 4, but now the underlying quality of the incumbent is not different from that of the alternatives $[\beta_{j=1,R}^{\hat{\beta}_R(0)} \equiv \beta_{j=2,U}^{\hat{\beta}_U(0)} \equiv \beta_i^{\hat{\beta}_i(0)}]$, and let $q_j \equiv \underline{q} < q_i \forall j$. Let q contribute to \bar{h} by a strictly, monotone decreasing and twice-differentiable function $v(q)$, so that $\underline{q} < q_i \forall j$ is sufficient to yield $\bar{h}_j > \bar{h}_i$ for $j = 1, 2$. Then the agent utilizing products according to the optimal dynamic allocation policy will always switch from the incumbent to the lower-priced alternative more quickly under the regulated product distribution.

Proof: Because the values q_i and q_j are known, static and factored into the values \bar{h} and $Q_i(\cdot)$, the agent will switch with to the alternatives with higher probability as $v(q_i) - v(q_j)$ increases. The

²⁰The quadratic form of the convergence ensues functionally from Ito’s Lemma in the present model. More broadly, it is a specific instance of the law of the iterated logarithm (Karatzas and Shreve 1991: 111-114).

difference in switching payoffs between regulated and unregulated products can be expressed as

$$\Delta\gamma^*(t) = \delta(Q_j(v(q_j)) - Q_i(v(q_i))) + \left(\frac{V^R(t)^2}{2\xi^2} \Psi_{\hat{\beta}\hat{\beta}}[\gamma(t), t] - \frac{V^U(t)^2}{2\xi^2} \Psi_{\hat{\beta}\hat{\beta}}[\gamma(t), t] \right)$$

By Lemmata 1 and 2, the rightmost quantity in brackets is strictly non-negative, hence $\gamma^{R*}(t) > \gamma^{U*}(t) \forall t$. QED.

The intuition of Proposition 6 is as follows; an agent in the unregulated market will hold on longer to an inferior (more costly) product given that her knowledge about the distribution of (potentially superior) alternative products facing her is subject to greater uncertainty and given that abandonment costs are strictly positive ($\chi_j > 0$). Price differentials are alone sufficient to generate switches, and these switches occur more quickly under approval-regulated product distributions than under distributions without approval regulation.

Intuitive Summary of the Bandit Problem and Regulation-Induced Utilization. The Agent’s learning problem, and its shaping by regulation, can be discussed verbally as follows. Imagine a countable set of products, and for each of these imagine an instantaneous draw from a normal distribution, which establishes the product’s initial estimate of quality (or ‘prior’) for the Agent’s learning problem. Without trying the product, this estimate is the only thing that the Agent knows about it.

Then the multi-armed bandit problem is just a product-by-product optimal stopping problem, but instead of the single barrier, there are a family of barriers (the “battery” of alternatives to the incumbent). Figure 1 displays this battery as several barriers, none of which touches the other (in finite time). If the evidence process for the incumbent product falls sufficiently low, the agent switches to the next-best product (the non-incumbent with the highest prior) and utilizes that next-best option, learning about it as an “experience good.” The optimal stopping problem is repeated, and the next-best product is used until it is abandoned in favor of the third-best product, and

so on. There is always positive probability that one product will be consumed forever because it delivers enough experienced value to avoid the switching barrier.²¹

Approval regulation has a number of effects upon this prior product distribution, but the salient one here is that it creates another (counterfactual) family of alternative products. This regulation-induced family is higher because $V_b^\Pi(t)^2$ is lower – there is less posterior variance (the Agent’s uncertainty) associated with the superior battery because regulation has reduced its prior variance (the uncertainty that the Agent would confront before utilizing or consuming any of the products).

6 Implications and Conclusions

Consistency with Studies of Approval Regimes. Interpreted accordingly, this result establishes a correspondence between the predictions of the current model and the interesting empirical results of Law (2003), Law and Kim (2005) and Law and Marks(2009). Law (2003) adduces panel regression evidence that state food regulation was associated with increased consumption of regulated foodstuffs. To be sure, state food regulation did *not* impose the approval-regulation scheme upon food products, but it did create quality-based entry barriers and in numerous states, laboratory and (animal) pharmacological testing of food additives, food preparation and processing processes and selected retail food products was a critical accompaniment to the regulatory institutions of the time. Along with the U.S. Department of Agriculture, such chemical and pharmacological testing of food and drug substances in the states was far progressed beyond other regulatory arrangements in the world or even in other realms of regulation (Law 2006). More relevant to the current model because licensure often involves a feature of pre-market testing, Law and Marks (2009) show benefits for minorities from professional licensure regimes because revealed information about worker quality can more easily trump discriminatory assumptions about minorities’ abilities or training. As they

²¹Importantly, consumption of this product forever is not a sufficient condition for showing that the product is truly the best among the countable alternative set; such behavior would be optimal only in the Bayesian and dynamic sense, given the Agent’s continuation values and given the Agent’s information at each stage of her decision process.

argue, “in those occupations in which minorities are underrepresented but for which information about worker quality is costly to obtain, licensing can reduce statistical discrimination” (2009: 364).

Policy Implications: Veto-Strong Regulators, Institutional Confidence, and Diffusion of Regulation-Induced Information. We have extended the analogy between safety regulation and financial regulation by focusing upon a critical property of institutions of safety regulation and a critical property of modern markets for complex products: (1) the regulatory veto and its induction of experiment, and (2) the central role of uncertainty in utilization decisions over time. Warren (2008) and others argue that product safety regulation ought to be a model for financial regulation in many forms. The argument here demonstrates that the critical mechanism that renders safety regulation powerful is the *forcible veto* (or removal power) of the government regulator that oversees the marketplace. It is this negative power that induces companies to provide more information about their products, and it is this negative power that induces firms to provide a particular kind of information (quasi-experimental information where the evidence is a truly random and representative draw from an underlying parameter of the product, with independent and identically distributed errors that permit valid inference). In the modern health world, the phased clinical studies system of the Food and Drug Administration (FDA) ratified in 1963 serve this important public function.²²

At its core, then, our model suggests that *some of the most important benefits of safety regulation have nothing to do with safety*. Approval regulation leads to a superior distribution of products, most likely, but regardless (and independent) of this result, approval regulation also leads to the provision of more information, and information of higher quality, than would be provided in the absence of approval regulation institutions. *Information provision alone is sufficient to induce*

²²The question of whether these rules and other forms of approval regulation are beneficial to society relative to their costs is of course an empirical proposition. Yet the model here shows that cost-benefit analyses of safety regulation that focus only upon “safety” benefits will generate systematically biased estimates of policy impact and will lead to invalid inferences (for one among many examples, see Peltzman (1974)). In other words, it is only when cost-benefit analyses examine the confidence effects of regulation that such exercises can even approximate the validity that public policy demands.

risk-neutral as well as risk-averse consumers to switch more readily to the products that will in fact benefit them.

Of course, our essay points to the crucial feature of public confidence in a regulator. In the model, if the Agent's prior variance is not reduced, then the confidence effects that generate more optimal switches among products do not materialize. If then the regulatory process is doubted, or the quality of the information produced by that process is subject to known bias. In recent years many regulators (particularly the FDA) appear to have suffered a material loss of public confidence. While a loss of regulatory confidence is not in and of itself sufficient to dash all of the benefits of approval regulation institutions (even poorly trusted regulators can prevent bad products from coming to market and thereby improve the quality of utilization), vital benefits of approval regulation are foregone when agents believe that the Regulator is being duped, is colluding with the Firm, or is otherwise negligent.

So too, it is critical that the experimentation conducted by the Firm and ratified by the Regulator be communicated transparently to consumers who can then access the information at low cost. This might occur through public disclosure of "stress test" data or data on various asset portfolios of investment holdings. In the model of this paper, this critical mechanism is accomplished by fiat; the cumulation of the product distribution from the approval regulation game is assumed to be observed by all, in a sort of "common knowledge" assumption common to decision theory and game theory. In practice, of course, even publicly disclosed experiments do not easily make their way into end-users' decisions among competing products, a fact which is well known in medicine and pharmaceuticals. Due in part to advertising but due as well to limited disclosure of product experimentation, the benefits of approval regulation are often poorly realized. This fact has led to plausible calls for "academic detailing" (where information provision is undertaken by disinterested (or less interested) parties; Avorn 2004) as well as public clinical trial registries. The public clinical trial registry model may be a useful information revelation device for those interested in "safety

regulation for financial products.”

Limitations and Extensions. As with all such theoretical efforts, our model has some basic limitations. What we have not yet done is to think about safety and the very real possibility that someone can get defrauded (the “exploding toaster” problem in Warren (2007, 2008)).²³ Such an extension is of vital importance, of course, but we have avoided it here for two reasons. First, the model demonstrates something that is not intuitive – that even when the Regulator and the Firm disagree only about unobserved quality, that information will be produced, and that this information will be valuable to a risk-neutral agent and will generate higher utilization of “good” products. Second, and less important, the particular mathematical solution concepts used here (infinitesimal generators and smooth pasting conditions) take quite different and less tractable forms in the presence of discontinuous data.

Relatedly, the model ignores not only “exploding toasters” but also human agents’ limited ability to discern between safe and unsafe toasters in a complex marketplace. Hence another important limitation of the present model is that it does not embed cognitive biases that are well known among students of psychology and its related fields. In the context of the multi-armed bandit problem modeled here, if human agents (consumers, end-users) rely upon biased information that they do not perceive as biased, then they will switch more readily to alternatives, but not necessarily to products that actually benefit them.

Another principal limitation of the model from the standpoint of mathematical generalizability is its separation of the regulation game and the consumer learning problem; there is no ‘full equilibrium’ analysis where the play in the regulation game is affected by anticipated reactions of a population of consumers. Another limitation comes in the simple and single-shot nature of the approval regulation game.²⁴ The main problem here is one of mathematical tractability. If a fully

²³See, however, Carpenter and Grimmer (2009) for a model where a regulator considers both continuous evidence of “quality” and discrete evidence of “safety” hazards.

²⁴But see Carpenter (2004) for a repeated optimal stopping model of the game where regulatory learning is enhanced in a repetitive context.

dynamic and strategic context were to be adopted, then the analysis would be vastly more complex and quite likely unresolvable analytically; at the very least, it would need to be conducted in the still-evolving and limited field of differential games.

Implications for Cost-Benefit Analysis, and Empirical and Historical Agendas. In addition, we think that critical research agendas lay in the empirical and historical extension of “confidence” effects in regulation. Historians and historically-oriented scholars may wish to ask a variety of questions: (1) how is trust created among wide and dispersed populations of human agents in response to consumers? (2) what are the historical mechanisms and contingencies by which regulation affects consumers’ beliefs about the set of alternative products they are facing? and (3) since there is no such thing as a single, undifferentiated consumer – instead there are heterogeneous audiences – how do differing perceptions of the regulator shape the confidence effects of institutions?

Empirical analyses, too, are needed, but our theory suggests that they should be focused not only upon the avoidance of drastically bad outcomes but on the quieter but no less vital outcome of everyday utilization of products. The empirical literature on regulation is shot through with attempts at weighing the safety- or externality-related “benefits” of the policy with its imposition-related “costs” (see the representative efforts of Viscusi (1993) or Boardman, Greenberg, Vining, and Weimer (1996)). Combined with the studies of Law (2003), Law and Kim (2005) and Law and Marks (2009), our model points to systematic patterns of regulatory effect that are not captured by this literature. At a minimum, many efforts in contemporary cost-benefit analysis of regulation are arguably affected by omitted variable bias. Consumer and public perceptions are far more than a matter of public opinion and survey research; they are integral to optimal administration of, and proper measurement of the benefits and costs of, a wide variety of regulatory policies.

Our near-term efforts (Carpenter, King, Law and Moffitt 2009) are focused upon an interesting and vital quasi-experiment that occurred in the area of pharmaceuticals, the FDA’s Drug Efficacy Study Initiative (DESI) of the late 1960s and early 1970s. Unlike previous interventions into the

therapeutic marketplace, the DESI was designed to remove patently safe products from the market because they had not proved their therapeutic efficacy by means of controlled studies (preferably randomized controlled studies with a placebo arm and with variable dosages). In other words, it was an act of efficacy regulation by a safety regulator; after DESI's completion, we hypothesize, physicians throughout the United States could prescribe more confidently knowing that the efficacy of most products had been supported by controlled clinical studies that were reviewed by a partially independent regulator.

Libertarian analyses of approval regulation – such as licensure, permitting, the governance of consumer products by the CPSC and the approval of new therapeutic products by agencies like the FDA – tend to predict a restriction of supply and a bevy of poor policy outcomes: a corresponding increase of commodity prices, reduction of product consumption, and in general a limitation of freedom. The theory presented here (and the set of related models can be built upon it) points to a quite different dynamic, one that is better attuned to the fact of learning (even imperfect and partially irrational learning) among human agents. Institutions of regulation that limit market entry but do so by compelling forms of experimentation or knowledge generation can often create a more predictable marketplace. In the marketplace created by approval regulation, both risk-neutral and risk-averse human agents will more readily enter this marketplace and more readily rely upon quality data to switch to the products that present them with the most value.

A Appendix

Proof of Lemma 1: Second-Order Stochastic Dominance of Regulation-Induced Product Distribution

Note that by the construction of the posterior variance for the Beta distribution $[= \frac{(\theta-\theta')(n+n'-\theta-\theta')}{((n+n')^2(n+n'+1))}]$, any additional experimentation will have a (quadratically) decreasing effect upon it. We first examine the properties of the ESE. For any first-period success there are two cases. When $k \leq \tilde{x}$,

$E_R[x]$ SOSD $E_U[x]$ because the high type would have submitted without experimentation. When $k \geq \tilde{x}$, $E_R[x]$ SOSD $E_U[x]$ because the Regulator's rejection strategy induces the low type to mix between submitting and experimenting. Because all high types would submit, the Regulator infers that the Firm is of the low type if he experiments. Thus a further experimental success ($e_2 = 1$) results in submission and acceptance, while a failure results in withdrawal because $k > m/(n+2)$. Since $\eta_F(1) > 0$ in the ESE under this condition, the experimentation (n') is strictly greater than under no regulation (where $\eta_F(1) = 0$).

Now suppose that costs are not high (so that EW is violated), and consider the path of a high type firm whose first experiment fails. The Firm cannot submit, but is sufficiently confident to conduct another experiment. With probability $m/(n+1)$, the Firm succeeds and then submits. But because the low type would also experiment with positive probability and submit if successful, the Regulator mixes between acceptance and rejection. Then $k \geq \tilde{x}$, $E_R[x]$ SOSD $E_U[x]$ because the high-type Firm would never have experimented in the first place. Second, suppose that $k > \tilde{x}$, the Firm is of the low type, and $e_1 = 1$. Because k is high, the Regulator does not accept all period 2 submissions and instead mixes. This makes the Firm indifferent between submission and further experimentation.

To complete the proof, note that by the posterior variance of the Beta distribution all late-submission equilibria will convey at least as much information for submitted/accepted products as will the early submission equilibria. **QED.**

Proof of Lemma 2

We begin by noting some facts about $G_U(x)$ and $G_R(x)$ which are drawn from Beta variates.

Fact 1: x and $E[x]$ are uniformly continuous on $[0, 1]$, and $E[x]$ and $Var[x]$ can be completely described by the space $\mathcal{C} = \mathcal{C}[0, 1]$ of continuous functions on the unit interval.

Fact 2: Any probability measures $P(x)$ and $P_n(x)$ describing x and $E[x]$ are tight in the sense that $\lim_{\epsilon \rightarrow 0} \sup_{|s-t| \leq \epsilon} |y(s) - y(t)| = 0$ for all $0 < \epsilon \leq 1$.

Convert $\pi_\zeta(x)$ into a series of mixtures with quadratically increasing additive terms, as follows.

$$[\text{1st mix}]: \pi_1(x) = (1 - \alpha_1)G_1(x) + \alpha_1G_2(x)$$

$$[\text{2nd mix}]: = \pi_2(x) = (1 - \alpha_1)((1 - \alpha_{1,1})G_{1,1}(x) + \alpha_{1,1}G_{1,2}(x)) + \alpha_1((1 - \alpha_{1,2})G_{2,1}(x) + \alpha_{1,2}G_{2,2}(x))$$

$$[\text{3rd mix}]: = \pi_3(x) = (1 - \alpha_1)\{(1 - \alpha_{1,1})[(1 - \alpha_{1,1,1})G_{1,1,1}(x) + \alpha_{1,1,1}G_{1,1,2}(x)] + \alpha_{1,1}[(1 - \alpha_{1,1,2})G_{1,2,1}(x) + \alpha_{1,1,2}G_{1,2,2}(x)]\} \\ + \alpha_1\{[(1 - \alpha_{1,2})[(1 - \alpha_{1,2,1})G_{2,1,1}(x) + \alpha_{1,2,1}G_{2,1,2}(x)] + \alpha_{1,2}[(1 - \alpha_{1,2,2})G_{2,2,1}(x) + \alpha_{1,2,2}G_{2,2,2}(x)]]\}$$

and so on, where again all f composing π_ζ are Beta-distributed.

Note that the ζ^{th} mix will have 2^ζ separate Beta distributions $G(\cdot)$ in the summand, and that distributions and their nearest mixture parameters (α) in the ζ^{th} mix will each be described by a ζ -dimensional index. This ζ -fold mix presents a flexible way of approximating successively refined divisions of “probabilities” on the space $\mathcal{C}[0, 1]$. By the *characteristic function lemma for mixture distributions* (Feller 1971: 504), each $\pi_\zeta(x)$ is itself a well-defined probability distribution on $[0, 1]$, hence if the G 's are well-defined, nondegenerate distributions then so are the π 's.²⁵

Now take the series $\frac{a'_{\zeta,1}}{b'_\zeta}, \dots, \frac{a'_{\zeta,i-1}}{b'_\zeta}, \frac{a'_{\zeta,i}}{b'_\zeta}, \dots, \frac{a'_{\zeta,N_\alpha}}{b'_\zeta}$ such that $\min_{1 < i < N_\alpha} \left[\frac{a'_{\zeta,i-1}}{b'_\zeta} - \frac{a'_{\zeta,i}}{b'_\zeta} \right] \geq \epsilon_\zeta$. Take further the series $0 = r_0 < r_1 < \dots < r_{N_\alpha} = 1$, arranged such that $\min_{1 < i < N_\alpha} (r_i - r_{i-1}) \geq \rho$. Then it is a consequence of the Arzela-Ascoli theorem (Theorem 7.4 of Billingsley 1999: 83) that

$$\Pr \left[y : 3 \max_{1 \leq i \leq N_f} \sup_{r_{i-1} \leq s \leq r_i} |y(s) - y(r_{i-1})| \leq 3\epsilon_\zeta \right] \leq \sum_{i=1}^{N_\alpha} \Pr \left[y : \sup_{r_{i-1} \leq s \leq r_i} |y(s) - y(r_{i-1})| \geq \epsilon_\zeta \right]$$

for arbitrary functions y on $\mathcal{C}[0, 1]$. Substitute $\frac{a'_{\zeta,i}}{b'_\zeta} - \frac{a'_{\zeta,i-1}}{b'_\zeta}$ for $y(s) - y(r_{i-1})$ and, fixing ϵ_ζ , let $a'_{\zeta,i}$ (along with $a'_{\zeta,i-1}$) pass with b'_ζ to infinity, passing from rationals to reals (Theorem 9.25 of Karatzas and Shreve 1989: 114, Billingsley 1999: 88-89). As $\left[\frac{a'_{\zeta,i-1}}{b'_\zeta} - \frac{a'_{\zeta,i}}{b'_\zeta} \right]$ gets small, so does the left-hand-side quantity in the Arzela-Ascoli result above. Then α is, by the modulus of continuity

²⁵The sizes of intervals α are given by the Poisson-Dirichlet distribution (Billingsley 1999: 39-43); this fact is very useful for applications of the framework here.

defined, a continuous probability distribution on $[0, 1]$.

Note that, because all G 's are Beta-distributed, so (by product results on Beta distributions; Jambunathan 1954; Kryszicki 1999) are all additive components of $\pi_\zeta(x)$. Hence each mix will now be a ζ^2 -term summation of Beta distributions, with mean and variance describable by continuous parameters. Normality of the integral distributions G^Σ follows from standard results on the sum of Beta distributions (Evans, Hastings and Peacock 1992: 36).

For preservation of stochastic dominance relations, note first that the product and summation operators in the integral distributions G^Σ preserve the continuity of countably additive functions of $\pi_\zeta(x)$, and G^R and G^U . Take any four distributions drawn from the exponential family (which include the Beta and normal), G_1 , G_2 , G_3 and G_4 . By additivity of independent variances, if G_1 SOSD G_2 , then for $0 \leq \alpha \leq 1$, $\alpha G_1 + (1 - \alpha)G_2$ SOSD $G_2 \equiv \alpha G_2 + (1 - \alpha)G_2$, and $\alpha G_1 + (1 - \alpha)G_3$ SOSD $\alpha G_2 + (1 - \alpha)G_3$. Then for the ζ^{th} and $(\zeta - 1)^{\text{th}}$ terms of any quadratic mixture as defined above, the following stochastic dominance relation holds:

$$[\zeta^{\text{th}} \text{ mix}]: \alpha_\zeta(\alpha_{\zeta-1}G_1 + (1 - \alpha_{\zeta-1})G_3) + (1 - \alpha_\zeta)(\alpha_{\zeta-1}G_3 + (1 - \alpha_{\zeta-1})G_4)$$

$$\text{SOSD } \alpha_\zeta(\alpha_{\zeta-1}G_2 + (1 - \alpha_{\zeta-1})G_3) + (1 - \alpha_\zeta)(\alpha_{\zeta-1}G_3 + (1 - \alpha_{\zeta-1})G_4)$$

Because ζ is arbitrary, $\sum_\zeta \left\{ \pi_\zeta^R(G^R)/N_\zeta \right\}$ SOSD $\sum_\zeta \left\{ \pi_\zeta^U(G^U)/N_\zeta \right\}$. But by Fatou's lemma,

$$\int_{\mathfrak{R}} \liminf_{\zeta \rightarrow \infty} \sum_\zeta \left\{ \pi_\zeta^R(G^R)/N_\zeta \right\} \leq \liminf_{\zeta \rightarrow \infty} \int_{\mathfrak{R}} \sum_\zeta \left\{ \pi_\zeta^U(G^U)/N_\zeta \right\}$$

Because second-order stochastic dominance simply invokes expectations of the variance of $G^{\Sigma,R}$ and $G^{\Sigma,U}$, then preservation of SOSD relations falls out as a special case for all finitely-valued G^Σ . The assumed square-integrability of G^Σ is sufficient, with Fatou's lemma, to preserve the finiteness of all expectations of G^Σ .

Lemma 3: Assume the hypotheses of Proposition 3, including the equivalence of histories $H[x_{1,t_1}] \equiv H[x_{2,t_2}]$. Now permit the equivalence of estimate histories to differ $\hat{\beta}_R[x_{1,t_1}] \neq \hat{\beta}_U[x_{2,t_2}]$,

in that

$$\hat{\beta}_{R,t} = \frac{\beta_R^0/v_b^R + x_1/\xi^2}{1/v_b^R + t_1/\xi^2} \text{ and } \hat{\beta}_{U,t} = \frac{\beta_U^0/v_b^U + x_2/\xi^2}{1/v_b^U + t_2/\xi^2}$$

Then the agent still switches more quickly to the regulated product (R) at $t_j = t_{\text{switch}}$ unless the prior difference is sufficiently large that

$$\begin{aligned} & \frac{(\hat{\beta}_U^0 - \hat{\beta}_R^0) v_U^{-1} v_R^{-1} + t_i \xi^{-2} (\hat{\beta}_U^0 v_U^{-1} + \hat{\beta}_R^0 v_R^{-1}) + \xi^{-2} (v_U^{-1} + v_R^{-1})}{v_U^{-1} v_R^{-1} + t_i \xi^{-2} (v_U^{-1} + v_R^{-1}) + t_i^2 \xi^{-4}} \\ & > \frac{\xi^2 + v_b t_{\text{switch}}}{\xi^2} \left[V_1^\Pi (t_{\text{switch}})^2 - V_2^\Pi (t_{\text{switch}})^2 \right] \end{aligned}$$

Proof of Lemma 3: The difference between the two estimates can now be expressed as

$$\frac{\hat{\beta}_U(0) v_U^{-1} + \xi^{-2} x_{2,t}}{v_U^{-1} + \xi^{-2} t_2} - \frac{\hat{\beta}_R(0) v_R^{-1} + \xi^{-2} x_{1,t}}{v_R^{-1} + \xi^{-2} t_2}$$

The difference in switching barriers is identical and unaffected by the differential drag on the coefficient histories. Solving the left-hand-side of the hypothesis for at least term scalar in $\hat{\beta}_U^0 - \hat{\beta}_R^0$ yields the solution. QED.

REFERENCES

- Alt, J., R. Calvert, and B. Humes. 1988. "Reputation and Hegemonic Stability: A Game-Theoretic Analysis." *American Political Science Review* 82: 445-466.
- Avorn, J. 2004. *Powerful Medicines: The Benefits, Risks, and Costs of Prescription Drugs*. New York: Knopf.
- Billingsley, P. 1999. *Convergence of Probability Measures*, Second Edition. New York: Wiley.
- Boardman, A. E., D. H. Greenberg, A. R. Vining, D. L. Weimer. 1996. *Cost-Benefit Analysis: Concepts and Practice*. Upper Saddle River, New Jersey: Prentice Hall.
- Calvert, R. 1987. "Reputation and Legislative Leadership." *Public Choice* 55: 81-119.

- Carpenter, D. P. 2004. "Protection without Capture: Dynamic Product Approval by a Politically Responsive, Learning Regulator." *American Political Science Review* 98(4): 613-631.
- Carpenter, D. P., and J. Grimmer. "The Downside of Deadlines," Robert Wood Johnson Scholars in Health Policy Working Paper #35, February 2009.
- Carpenter, D. P., J. King, M. Law and S. Moffitt. 2009. "Therapeutic and Economic Effects of Efficacy-Based Drug Withdrawals: The Case of the Drug Efficacy Study Initiative," working paper [in progress], Department of Government, Harvard University, 2009.
- Carpenter, D. P. and M. M. Ting. 2007. "Regulatory Errors with Endogenous Agendas," *American Journal of Political Science* 51 (4) (October) 835-853.
- Chernoff, H. 1968. Optimal Stochastic Control. *Sankhya The Indian Journal of Statistics Series A* 30 (3): 221-252.
- Dixit, A., and R. Pindyck. 1994. *Investment Under Uncertainty*. Princeton: Princeton University Press.
- Evans, M., N. Hastings, and B. Peacock. 1993. *Statistical Distributions*, Second Edition. New York: Wiley.
- Feller, W. 1971. *An Introduction to Probability Theory and Its Applications*, Volume II. New York: Wiley.
- Heimann, C. F. L. 1997. *Acceptable Risks: Politics, Policy, and Risky Technologies*. Ann Arbor: University of Michigan Press.
- Hilts, P. 2003. *Protecting America's Health: The FDA, Business, and One Hundred Years of Regulation*. New York: Knopf.
- Jambunathan, M. V. 1954. "Some Properties of Beta and Gamma Distributions," *Annals of Mathematical Statistics* 25: 401-5.
- Karatzas, I. 1984. "Gittins Indices in the Dynamic Allocation Problem for Diffusion Processes," *The Annals of Probability* 12 (1) 173-192.
- Karatzas, I., and S. Shreve. 1991. *Brownian Motion and Stochastic Calculus*. Second Edition. New York: Springer-Verlag.
- Krysicki, W. 1999. "On some new properties of the Beta distribution," *Statistics and Probability Letters* 42 (1999) 131-137.
- Landau, M. 1969. "Redundancy, Rationality, and the Problem of Duplication and Overlap." *Public Administration Review* 29(4): 346-358.
- Law, M. 2003. "The Origins of State Pure Food Regulation." *Journal of Economic History* 63 (4): 1103-1130.
- Law, M. 2006. "How do Regulators Regulate? Enforcement of the Pure Food and Drugs Act, 1907-38." *Journal of Law, Economics, and Organization* 22 (2): 459-89.
- Law, M., and S. Kim. 2005. "Specialization and Regulation: The Rise of Professionals and the Emergence of Occupational Licensing Regulation," *Journal of Economic History* 65 (3): 723-756.

- Law, M., and M. S. Marks. 2009. "The Effects of Occupational Licensing Laws on Minorities: Evidence from the Progressive Era" *Journal of Law and Economics* 52 (2): 351-66.
- Mandelbaum, A. 1987. "Continuous Multi-Armed Bandits and Multiparameter Processes," *The Annals of Probability* 15 (4) 1527-1556.
- Moscarini, G., and L. Smith. 2001. "The Optimal Level of Experimentation." *Econometrica* 69(6): 1629-1644.
- Peltzman, S. 1976. *Regulation of Pharmaceutical Innovation: The 1962 Amendments*. Washington, D.C.: American Enterprise Institute.
- Protter, P. E. 2005. *Stochastic Integration and Differential Equations*. New York: Springer.
- Quirk, P. J. 1980. "Food and Drug Administration," in J. Q. Wilson (ed.), *The Politics of Regulation*. New York: Basic Books.
- Reinganum, J. F. 1982. "A Dynamic Game of R and D: Patent Protection and Competitive Behavior." *Econometrica* 50(3): 671-688.
- Shiryaev, A. N. 1978. *Optimal Stopping Rules*. New York: Springer-Verlag.
- Simon, H. A. 1968. *Administrative Behavior*. New York: Free Press.
- Stiglitz, J. 2008. "The Financial Crisis of 2007/2008 and its Macroeconomic Consequences. Working paper, Columbia University.
- Ting, M. M. 2003. "A Strategic Theory of Bureaucratic Redundancy." *American Journal of Political Science* 47(2): 274-292.
- Viscusi, W. K. 1993. "The Value of Risks to Life and Health," *Journal of Economic Literature*, 31 (4): 1912-1946.
- Warren, E. 2007. "Unsafe at Any Rate," *Democracy: A Journal of Ideas* Issue #5, Summer 2007.
- Warren, E. 2008. "Product Safety Regulation as a Model for Financial Services Regulation," *The Journal of Consumer Affairs*, 42 (3) 452-60.