Protection without Capture: Product Approval by a Politically Responsive, Learning Regulator

DANIEL P. CARPENTER *Harvard University*

When policy arrangements appear to favor well-organized and wealthy interests, should we infer "capture" of the political process? In particular, might larger firms receive regulatory "protection" even when the regulatory agency is not captured by producers? I model regulatory approval—product approval, licensing, permitting and grant making—as a repeated optimal stopping problem faced by a learning regulator subject to variable political pressure. The model is general but stylistically applied to pharmaceutical regulation. Under the assumption that consumers are differentially organized, but producers are not, there nonetheless exist two forms of "protection" for larger, older producers. First, firms submitting more applications may expect quicker and more likely approvals, even in cases where their reputations for safety are below industry average. Second, "early entrants" to an exclusive market niche (disease) receive shorter expected approval times than later entrants, even when later entrants offer known quality improvements. The findings extend to cases of bounded rationality and a reduced form of endogenous firm submissions. The model shows that even interest-neutral "consumer" regulation can generate protectionist outcomes, and that commonly adduced evidence for capture is often observationally equivalent to evidence for other models of regulation.

> *The theory [of economic regulation] tells us to look, as precisely and carefully as we can, at who gains and who loses, and how much, when we seek to explain a regulatory policy . . . . It is of course true that the theory would be contradicted if, for a given regulatory policy, we found the group with larger benefits and lower costs of political action being dominated by another group with lesser benefits and higher cost of political action . . . .*
>
> *The first purpose of the empirical studies is to identify the purpose of the legislation! The announced goals of a policy are sometimes unrelated or perversely related to its actual effects, and* the truly intended effects should be deduced from the actual effects.
>
> —George J. Stigler, "The Theory of Economic Regulation" (1975, 140)

> *The prospect for any change in the [OSHA cotton dust] standard, however, is not great. Now that the large firms in the industry are in compliance, they no longer advocate changes in the regulation. Presumably, the reason is that the capital costs of achieving compliance represent a barrier to the entry of newcomers into the industry. This is simply one more illustration of the familiar point that surviving firms often have a strong vested interest in the continuation of a regulatory system.*
>
> —W. Kip Viscusi (1992, 177)

Does government regulation favor larger, older producers and impede smaller, newer firms? If so, then why does government regulation have such disparate effects? For decades, political scientists and economists have turned to the "capture" theory of regulation to explain these established-firm advantages. Capture theory posits that larger and older firms use regulation as a political substitute for economic competition, constructing entry barriers against their smaller and newer competitors (existing or potential) or using regulation to impose disproportionate costs upon smaller and newer firms (as Viscusi [1992] suggests). Because larger firms are better able to organize politically, the argument goes, they can induce politicians and bureaucrats to behave in ways that constrain smaller firms. Stigler's moral is lucid: Given observed firm advantages under a given regulation, *we ought to infer producer capture* regardless of the stated purposes of the law.

The reach of the capture argument is substantial; it concerns general principles of politics and public policy. Under what conditions—and with what evidence—should policymakers, scholars, and citizens conclude that a government agency or program is politically influenced or captured? When we observe that the "actual effects" of policy seem to favor one group or firm over another, should we proceed as Stigler recommends and deduce that capture or some form of distributive politics is at work? The central thrust of this essay is that in many scenarios we should not

Daniel P. Carpenter is Professor, Department of Government, Harvard University, 1805 Cambridge Street, Cambridge, MA 02138 (*dcarpenter@latte.harvard.edu*).

For helpful comments and discussions, I thank three anonymous reviewers, the Editor, John Brehm, Brandice Canes-Wrone, Scott Gates, Rick Hall, Thomas Hammond, Robert Lieberman, David Meltzer, John Moran, Mary Olson, Ariel Pakes, Mark J. Schlesinger, Liam Schwartz, Ken Shotts, Dean Smith, Michael Ting, Greg Wawro, Natasha Zharinova, and audiences at the Michigan State Methods Workshop, the American Politics Workshops at Columbia and Harvard, the Yale Public Policy Workshop, and the Texas A&M Conference on the Scientific Study of Bureaucracy. I acknowledge the National Science Foundation (SES-0076452), the Robert Wood Johnson Foundation Scholars in Health Policy Program, the Harvard University Department of Government, and the Center for Advanced Study in the Behavioral Sciences for research support. I am solely responsible for all remaining errors and interpretations.

do so, because observed patterns of policy advantage and regulatory advantage are an insufficient condition for a valid inference of capture. Such patterns may exist for many reasons unrelated to distributive politics, and may exist even when the policy is neutrally implemented. The pattern of inference recommended by Stigler (1975; and widely used by Viscusi and many others) is, then, often unwarranted.

In this essay, I reconsider firm-based disparities in government regulation and ask whether they can prevail even when regulation is not "acquired by the industry," as Stigler would posit. I ponder mathematically the world in which "capture" does not exist but protection still might. Consider, then, the following thought experiment: Suppose that a regulatory process were established in which all firms were identical in terms of political influence and quality. In this world, no firm would possess either (1) political clout or (2) a quality advantage over its competitors. Would larger and older firms still receive more favorable decisions from government regulators? The theory presented here answers in the affirmative. I offer an optimal stopping model of product approval that formalizes several noncapture mechanisms that may account for larger- and older-firm advantages in regulation. These mechanisms may cloud inferences of the sort that Stigler recommends, but they may also *interact* with capture to boost the advantages of larger firms. I focus here on the duration of regulatory decisions, with particular applications to pharmaceutical regulation, an exemplary case in which larger-firm advantage has been found (Grabowski and Vernon 1976; Olson 1997; Thomas 1990).

Delay, Familiarity, and Consumer Pressure: Why a Neutral Regulator Might Favor Established Firms

When regulators face a stopping problem, or a decision of *when to take costly action that is also costly to reverse*—as in product approval decisions (Heimann 1997), licensing (Spence 1999), or even enforcement litigation (Gordon 1999)—larger and older firms have several advantages. First, firms that market more products will be better known to the regulator, who will have less uncertainty over the firm's underlying qualities, for instance, the ability to produce a safe product. I show that even when regulators are neutral with regard to risk, reduced uncertainty about firm attributes leads to quicker decisions. Second, larger firms with greater capitalization will often enter given market niches earlier, and these niches may contain organized consumers (e.g., AIDS sufferers) that the regulator wishes to satisfy. Where this pressure is strong, such "early-entrant protection" implicitly benefits large firms. Finally, because dynamically learning regulators are unlikely to make immediate approval decisions, larger firms usually benefit because time is less costly for them relative to their assets. In the case of product approval, a delay for one product may cripple a start-up firm, whereas a more established company can more easily absorb the costs of delayed entry such as reduced cash flow, prolonged investor uncertainty, and the like.

These advantages are shown to be quite robust. In particular, the advantage of familiarity holds even in cases where the familiar firm has a bad reputation for product safety. In addition, early-entrant advantages exist even under cases in which later entrants—products that enter the market for a given niche well after the "early entrants" do—offer superior products.

For two reasons, my purpose here is *not* to offer a fundamental critique of the capture theory of regulation. That has been done elsewhere, for one (Breyer 1982; Wilson 1980). A better reason is that the present model actually illuminates the study of rent-seeking in regulatory politics. As the model shows, there may be complementarities between capture-based factors and noncapture factors in effecting larger-firm advantages. In industries where larger firms possess political rents, these advantages may be heightened by the dynamics illuminated here.

THE LOGIC OF LARGE-FIRM ADVANTAGE IN REGULATION

"Larger" or "established" firms in an industry differ from their "smaller" and "newer" counterparts by greater revenues, more products, greater diversification of product lines, larger capitalization (more assets), and enhanced tolerance of financial exposure. They also enjoy relative advantages under numerous forms of regulation—ranging from price-setting institutions to content regulation (e.g., the Federal Communications Commission) to "consumer safety" measures. Such advantages appear in numerous industries, ranging from transportation (Rothenberg 1994) to telecommunications (Crandall and Flamm 1989; Noll 1973) to pharmaceuticals (Grabowski and Vernon 1983).

In few markets is the power of the state to reward firms greater than this last one, the pharmaceutical industry. Throughout the industrialized democracies, particularly in the United States, the very legal marketability of a drug product depends on *prior* approval by government regulators. As a result of the 1938 Food Drug and Cosmetic Act, the U.S. Food and Drug Administration (FDA) has sole authority to approve prescription drugs for marketing. The state remains the ultimate gatekeeper in the pharmaceutical industry, both here and overseas.

The existence of systematic regulatory advantages for larger firms in the pharmaceutical industry has been well documented. Grabowski and Vernon (1976) argue that "upward shifts in costs and risk produced by increased [pharmaceutical] regulation . . . operate to concentrate innovation in fewer and larger firms" (190). Thomas (1990) finds that the 1962 Amendments to the Food, Drug and Cosmetic Act—which required drug companies to show *effectiveness* in addition to safety before marketing—caused heavier reductions in product innovation among small firms than among large firms. More recently, Olson (1997) finds quicker approval times for firms with greater employment and

more FDA product applications, and Carpenter (2002) and Kyle (2002) find that the speed and probability of regulated entry into pharmaceutical markets are (1) an increasing function of firm market experience and (2) a decreasing function of previous entrants. Employment and experience are characteristics highly correlated with firm size, and early entry into a market is *per se* correlated with age. Perhaps most crucial to these advantages is the finding that larger and older firms receive quicker approvals for their new drug applications (NDAs). Whatever advantages large firms have in product development seem to be magnified by FDA *decision making*.

The capture theory formulated by Bernstein (1955) and Huntington (1953) and rendered more explicit by Peltzman (1976) and Stigler (1975) offers a simple and powerful explanation for these advantages. Because larger firms have lower marginal costs and higher marginal benefits of political action, they can more easily capture the regulatory process, either by shifting legislation to their advantage (Stigler 1975, 123–28) or by shifting the administration of an otherwise costly law to their advantage (Stigler 1975, 162–66).[1]

In the past two decades, scholars have refined the capture perspective by advancing "interest-group" and "rent-seeking" theories of regulation (Becker 1985; Rothenberg 1994). These views concede that producer interests are not monolithic—firm interests vary by specialization and other things—and that consumers can also organize. Yet even these theories rarely shy away from the premise that large firms in an industry tend to dominate regulation. As Rothenberg (1994, 4) argues, "The view that producer dominance is the modal description of regulatory politics remains perhaps the primary means of conceptualizing them for popular commentators and scholars alike." And as the opening epigraph from Viscusi (1992) suggests, scholars are still apt to follow Stigler's advice for proper inference when large-firm advantage is observed (see also Bartel and Thomas 1987). For purposes of this paper, then, the proposition that lies at the heart of a "capture" account of regulation is that, by legislative structure or by cozy relations with the industry, regulators systematically favor large, established firms over newer entrants and smaller firms.

Stigler's argument is a powerful one and has usefully aided a generation of regulation scholars. Yet its counsel that regulatory *results* justify inference about regulatory *mechanisms* leaves open the possibility that numerous factors may explain observed regulatory outcomes. In particular, there may be cases where producers have not "captured" a legislature or agency but in which large-firm advantage prevails nonetheless.

Scholars have begun to recognize this possibility. Olson (1997) has employed "external signals" theory (Joskow 1974; Noll 1985) to suggest that regulators see firm experience as one of several "signals" of likely product quality—other signals being R&D and productivity—and so condition their reviews accordingly. Olson finds systematic evidence for this hypothesis in a regression analysis of drugs approved from 1990 to 1992. Controlling for clinical factors and FDA priority ratings, Olson finds that firms with greater R&D, more employees and more FDA product applications receive shorter approval times for their drug applications. A 1995 study by the General Accounting Office (GAO; 1995, 32) also finds that firms with more submissions receive shorter review times. Employment and submissions are highly correlated with firm size. Olson construes these results to mean not that larger firms capture drug regulation but, rather, that a rational agency interprets size-related firm characteristics as signals or correlates of unobserved product quality.

If Olson's argument is correct, then larger-firm advantage can exist even when regulators seek only to reduce their Type I and Type II errors. Yet is this result generalizable? Would mathematical formalization support the argument that an uncertain regulator will indeed act more quickly when firm characteristics shed light on product applications? What if the regulator is risk-neutral? and What if the agency is in fact influenced by the political clout of large firms? I am unaware of any attempt to construct a systematic formal theory that explores the firm-specific implications of approval regulation by a politically responsive, Bayesian agency.

A DYNAMIC MODEL OF APPROVAL REGULATION

Product approval presents the regulator with a learning problem. The regulator must review a new product application (with accompanying data) and decide when the apparent benefits of the product outweigh the costs or risks associated with its use.

Reputation and the Value of Delay

The basis of the model is that product approval regulators guard their reputation for protecting consumer safety. In the case of the FDA, this reputation may or may not be "deserved" in the sense that historically observed patterns of drug safety may be attributable not to regulation but to quality improvements by firms. Nonetheless, I assume that such a reputation exists. Observers of the 1962 Amendments argue that the FDA's decision not to approve the drug thalidomide played a crucial role in the expanded discretion that the agency received in the new law (Hilts 2003). Historians (Jackson 1970) and political scientists (Quirk 1980) have noted the political asset value of this reputation. It can be used to generate public support, to

[1] References to Stigler's original (1971) essay refer throughout to the page numbers in *The Citizen and the State*. Stigler argues mainly in terms of the *industry* erecting entry barriers to potential competitors, but his argument is easily rendered in terms of larger firms erecting barriers to smaller firms *within* an industry. To begin with, Stigler frequently uses "the industry" to denote the larger and established firms in a market. Moreover, if "smaller" firms are firms operating on the boundaries of the industry, or as firms that were formerly potential entrants, then the entire theory elaborated in Stigler 1975 and Peltzman 1976 applies straightforwardly to large- and older-firm advantage. More recent analysts such as Viscusi (1992), Bartel and Thomas (1987) and Thomas (1990) interpret capture in terms of large-firm regulatory advantage.

achieve delegated authority and discretion from politicians, to protect the agency from political attack, and to recruit and retain valued employees (Carpenter 2001).

Given this concern for reputation, the agency views its approval decisions as fundamentally irreversible (or reversible only at cost). The fundamental trade-off facing the regulator is one of time. Early approvals benefit both the producing firm and the organized consumers who may demand new products, and these interests can make it politically costly for the agency to delay. Yet regulatory delay buys time and information—the ability to review the firm's potential new product more carefully and to request more studies. As in financial options, there is an informational value to waiting, which is marginally decreasing as the agency learns more (see Ting 2003 for a similar but more general model of bureaucratic learning).[2]

The problem facing the regulator is one of stopping the review only once the payoff of approval exceeds *both* the reputational losses associated with the danger of the product *and* the value of waiting for more information. In other words, simple cost–benefit analysis applies poorly to this important class of administrative choices. The rational agency does *not* approve the product when its apparent danger is less than the payoff of approval but will wait further.

Assumptions and Parameters

I stylize the model for the drug approval process as it operates in the United States and other countries. By regarding diseases as "markets" or "market niches," curing probabilities as product quality, and danger as general "hazards" of a product, the model can be applied to other situations. Let all products be indexed by i, markets (diseases) by j, and firms by k. The model assumes an exogenous industry production process in which the agency expects products to be submitted at a constant rate $\lambda_j > 0$ over time. (I discuss this under The Problem of Endogenous Submissions, below.) Products are assumed to have "niche specificity"; that is, drugs can treat one disease only.[3]

All products in the model are characterized by two parameters. First, let γ_{ij} $(0 \leq \gamma_{ij} \leq 1)$ be *product quality* (for drugs, a *curing probability* of the drug, which can be interpreted as the fraction of people with disease j that drug i will cure). We assume that γ_{ij} is fixed and known with certainty throughout the agency's decision problem.[4]

Second, let μ_{ij} be the *danger* of the product, which can be thought of as the expected number of people that will be harmed or killed by the product over a given interval of time. Normalizing the interval to one, μ_{ij} may be thought of as the *rate* of harming consumers. The greater the danger of the product, the more its approval will harm the agency's reputation for protecting public safety.

I assume throughout that a product's danger is independent of its quality, which implies $\text{cov}(\mu_{ij}, \gamma_{ij}) = 0$. Because the product quality γ_{ij} is niche-specific, this independence is understood in a particular way. It is possible that a "good" cancer drug could cure a patient's cancer but induce such severe "side effects" in doing so that the patient dies of other causes (e.g., liver toxicity). I assume danger and quality are separable because quality is disease-specific (the cure of a specific ailment), whereas danger refers to literal "*side* effects," namely, the harm a drug may cause to some organ or physiological process in the person's body *other* than that which was the drug's intended target.[5]

Learning about Danger in Continuous Time. The agency observes a series of experiments (e.g., clinical trials) in which a product either harms or does not harm the consumer. The sequence of binary outcomes—"harm" or "not harm"—becomes the "data" for the agency's decision. The model posits the evolution of these observations as a continuous-time Wiener process, or "Brownian motion."[6] Casually, we may think of Brownian motion as an "all-purpose" random process in continuous time whose movements are described by the Normal distribution.

Observed harm in regulatory review evolves according to a Wiener process $X_{it} = X(t)$, a linear function of underlying danger (μ_{ij}) plus a random component, or

$$X(t) = \mu_{ij}\, t + \sigma_{ij}\, z(t), \qquad \textbf{(1)}$$

where μ_{ij} and σ_{ij} are constants and $\sigma_{ij} > 0$, and where $z(t)$ is a standard normal variable with mean zero and variance t. A more "dangerous" product—one with higher μ—will cause more harm, but harm will also be influenced by the random term $z(t)$. A higher σ yields a more volatile review, namely, a series of experiments from which the agency will find it harder to learn what μ is.

Danger as the Cost of Approval. Observed harm is a Markov process, and the agency can learn about μ

[2] The agency as I model it is always averse to "danger" (more hazardous products) but is also "risk-neutral." Intuitively, risk-neutrality implies that, given an expected outcome, the regulator is indifferent to higher or lower levels of variance around that outcome. Put differently, the agency's concern about uncertainty is endogenous to this model, not assumed.

[3] The one drug- one disease assumption can be relaxed without affecting the general results of the model. Clearly the exogeneity assumption here is violated in practice. Where the agency becomes more stringent in its product approval decisions, firms will develop fewer drugs and do so more slowly (Grabowski and Vernon 1983; Peltzman 1973; Thomas 1990). For a model that endogenizes submissions, see Carpenter and Ting 2004.

[4] One may also interpret γ_{ij} as the equilibrium market share upon the product's approval. The agency considers both safety and efficacy in this model, but for simplicity, I assume that only danger is learned and subject to uncertainty. The model can be adapted to the case where γ_{ij} is also learned, but the dynamics become quite complex when the agency updates more than two processes, and the simpler variant here adequately describes the review process.

[5] I thank an anonymous reviewer for this suggestion and for the clarifying language.

[6] As Dixit (1993, 2–3) shows, the Wiener-process representation can also be derived by taking limits toward zero on a discrete-time walk of binary outcomes.

in a simple Bayesian fashion based on the observed history of $X(t)$. Letting $X(t)$ start arbitrarily at 0, then it is normally distributed with mean μt and variance $\sigma^2 t$. We assume that σ is the same across products but that μ differs across them, according to a normal distribution with mean m and variance s. For any product review of length t and accumulated harm $X(t) = x$, the dual $[x, t]$ is a sufficient statistic for the history. In other words, the agency knowing the entire clinical history is no better off than if only $[x, t]$ were revealed. Given these sufficient statistics, Bayesian estimates of μ are

$$\text{posterior mean} \equiv E_{xt}(\mu_{ij}) = \hat{\mu}_{ijt} = \frac{(m/s) + \left(x/\sigma_{ij}^2\right)}{(1/s) + \left(t/\sigma_{ij}^2\right)}, \tag{2a}$$

$$\text{posterior variance} \equiv S(t) = \frac{1}{(1/s) + \left(t/\sigma_{ij}^2\right)}. \tag{2b}$$

Notice that

$$\lim_{t\to\infty} \hat{\mu}_{ijt} = \frac{x}{t} = \mu_{ij} \quad \text{and} \quad \lim_{t\to\infty} S(t) = 0. \tag{3}$$

The posterior variance $S(t)$ is the agency's expected uncertainty about the true value of μ. For this reason, the value of delaying another moment (another dt) is an increasing function of $S(t)$.[7]

The agency's goal during product review is to estimate the danger of the product (μ). Although real-world agencies have the option to recall a bad product (or induce producers to recall it), the agency has no such option here. The reason concerns reputation protection. Once a product has done sufficient harm to warrant a recall, the agency cannot recover its reputational losses via recall because observers will infer that the agency has made a "bad" decision. Hence the decision to approve a product is reputationally irreversible.[8] The political cost of approval, then, is the reputational loss that accrues to the agency from the danger of the product. Upon product approval, then, the agency loses μ, which can be learned only through preapproval review.

The Political Demand for Products and the Approval Payoff. One way for the agency to protect its reputation would be to review submissions forever (or reject them all). Yet there is often a *political demand* for products makes this delay costly. Drug approval provides a lucid example. In recent years the FDA has appeared to be highly responsive to the demands of (potential) drug consumers (Olson 1995; Vogel 1996). Firms also attempt to place pressure on the agency for quick approvals. Organized consumers and producers lobby the agency directly but also apply pressure indirectly through elected politicians.

For this section I define the agency's *approval payoff* from approving the N_j-th product as a function only of market attributes such as market size (disease prevalence), intensity of demand or "willingness to pay" (disease severity), public salience, and the political organization of consumers. The payoff may be written

$$A = f(\psi_j, L_j, N_j), \tag{4}$$

Here L_j is the potential size of the market (in pharmaceutical settings, disease j's *prevalence*, or the number of persons with disease j).[9] ψ_j is the *political multiplier* of market j, a positive parameter.[10] ψ_j can be interpreted as the expected number of citizens, for every citizen in market j, who will apply pressure on the agency or the politicians governing it. (In pharmaceutical settings, this is the number of citizens per patient with disease j. In some [but not all] cases, this parameter may capture the political impact of disease severity). N_j is the number of marketed products that have already entered market j (in a pharmaceutical context, the number of available drugs that already treat disease j).

I postpone exact specification of A to the section on early-entrant protection. For now, a crucial aspect of equation (4) is that none of its constituent variables is indexed by k. That is, I assume (for now) that the approval payoff is unaffected by firm characteristics such as contributions, reputation, and rents. In other words, the agency is not "captured" in the sense that its preferences do not direct it to treat some firms differently from others per se.

The Agency's Optimal Policy. The problem facing the agency can be described as the optimal stopping of the process $\hat{\mu}_t$, with the following objective (suppressing some subscripts):

$$\begin{aligned} &\max E_e^{-\delta(t_{app})} \left\{ A - E_{\hat{\mu},t} \int_t^{\infty} e^{-\delta(y-t)} \mu^*(y, \omega)\, dy \right\} \\ &\quad = E e^{-\delta(t_{app})} \left\{ A - \delta^{-1} \mu^*[t_{app}, \omega] \right\}, \end{aligned} \tag{5}$$

where δ is the discount factor, t_{app} is a given approval time, μ^* is the agency's estimate of danger at the optimal stopping time (as given in Eqs. [2]), ω denotes an elementary event in the probability space Ω, and y is a variable of integration.

The regulator's optimal policy is to divide the space of possible outcomes into two regions (DeGroot 1970)—a *continuation region*, where observed values

[7] See Chernoff (1968) and later, Jovanovic (1979), for a similar modeling architecture. I rely throughout on the scale-invariance of Brownian diffusions (Miroschnichenko 1975, 389–91).

[8] Other functional forms are possible. If the agency's reputation has the character of a durable asset subject to erosion, then perhaps the regulator's utility is dependent not on μ, but on $\exp(\mu)$. This would render the stochastic process a "geometric" Brownian motion (Dixit 1993), but none of the substantive results of the model differ.

[9] The constancy of L_j is a crucial assumption for the solution concept of smooth pasting. For now, simply assume that a rapidly growing disease has a high political multiplier. I also assume that indexation across j offers no information about danger; clinical risk and indications are assumed independent. In practice, this is almost certainly violated.

[10] The parameter can exceed one because citizens other than those directly afflicted by a disease—relatives, friends, or other allies in organized patient associations and the media—may lobby for drug approval.

FIGURE 1. Simulated First Passage of the Danger Process $\hat{\mu}(\mu, t)$ through the Agency's Approval Barrier $\eta(t)$

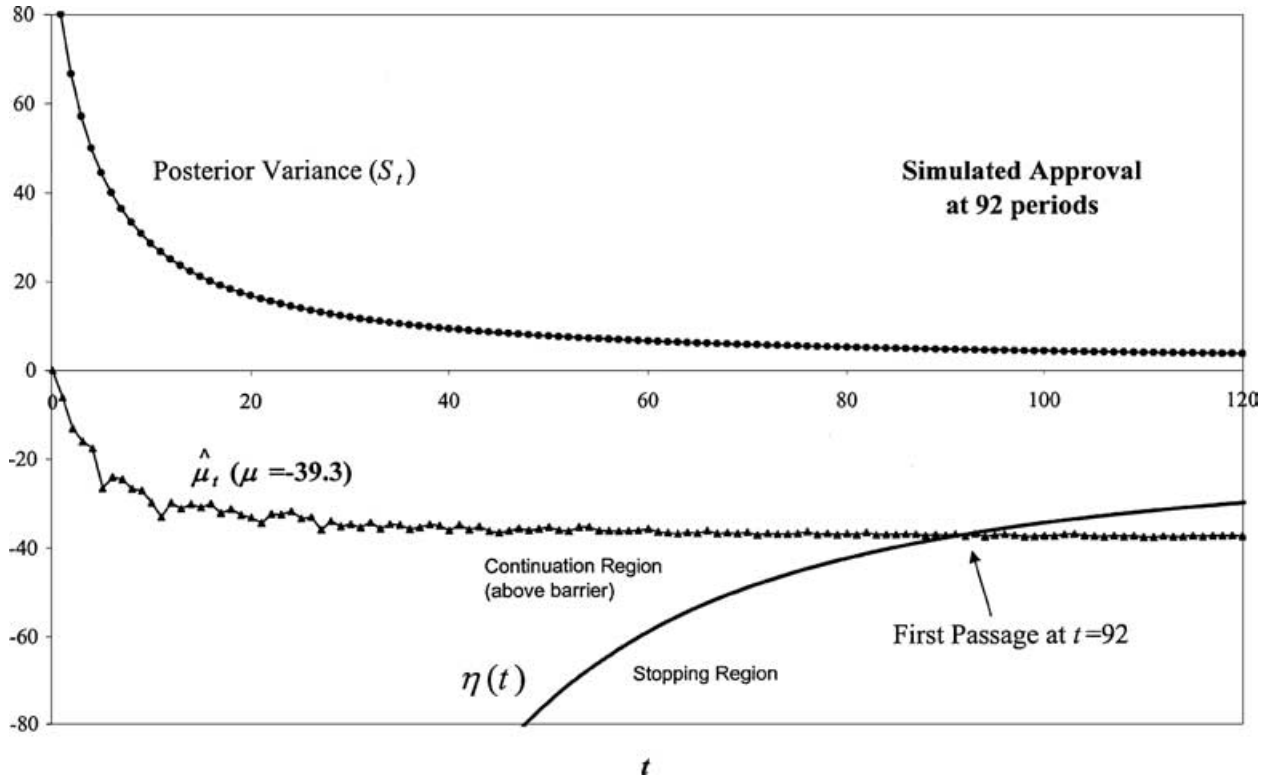

of estimated danger $[\hat{\mu}_t]$ suggest waiting, and a *termination region*, where values of danger suggest "stopping" the review process and approval of the product. Figure 1 offers a sample division of the space. Proposition 1 characterizes the agency's best policy.

Proposition 1. *For the approval problem characterized by* (5)*, there exists an optimal stopping policy characterized by the unique partitioning of* $[\hat{\mu}_t, t]$ *into a continuation region C and a termination region T, such that the approval barrier* $\eta^*(t)$ *represents the boundary of C, and the first passage of* $\hat{\mu}_t$ *through* $\eta^*(t)$ *represents the optimal approval of the product, where* $\eta^*(t)$ *is*

$$\eta^*(t) = \delta A - \frac{S(t)^2}{2\sigma^2} F_{\hat{\mu}\hat{\mu}}[\eta(t), t]. \qquad \textbf{(6)}$$

Proof. Proofs of all propositions, lemmata, and corollaries are given in the Appendix. ■

By Eq. (3), the limit of $\eta^*(t)$ is δA, which $\hat{\mu}_t$ approaches from above. The fully discounted benefits of approval (which have remained constant throughout the decision) are equal to δA. Meanwhile, the cost of approval is μ. In the limit, the estimated danger $\hat{\mu}_t$ converges to the true danger μ, so that the product is approved asymptotically if $\mu < \delta A$. In finite time, of course, the agency cannot make the decision in such terms. There is always a value to waiting for more information, regardless of the agency's attitude toward risk.

Characteristics of the Approval Distribution. Let $G^*(t)$ be the approval distribution, or the probability of approval at time t under the optimal policy. Proposition 2 notes an interesting property of $G^*(t)$: Even when the agency cannot formally reject a product, some products will still be "rejected" in the sense that they will never be approved.

Proposition 2. De Facto Rejection without a Rejection Option. *In expectation, a nonzero fraction of products will never receive the agency's approval, as* $\lim_{t\to\infty} G^*(t) < 1$.

In two ways, Proposition 2 coheres with actual regulatory behavior, for instance, the FDA. It is trivially consistent, of course, with the fact that not all drug submissions get approved. Yet it is also consistent with the fact that *the FDA never formally rejects a drug*. It simply deems a drug "not approvable," and nothing prevents the producing company from submitting more data about the drug if it seeks approval at a later date (GAO 1995).

Comparative statics are based on the following result.

Proposition 3. *Let* $E[t_{app} \mid t^*_{app} < \infty]$ *be the expected approval time under the optimal policy. Then*

conditioned on approval, $E[t_{app} \mid t^*_{app} < \infty]$ is strictly decreasing in the payoff A.

From Proposition 3, any variable in which the approval payoff is increasing is also a variable in which expected approval times are *decreasing*. Another prediction of the learning model concerns the speed of approval.

Proposition 4. The Scarcity of Quick Approval. *Let $g^*(t)$ be the density of $G^*(t)$. For any product $i[t_{app}(i) < \infty]$ such that an optimal approval time exists, the approval hazard $\theta(t) = g^*(t)/[1 - G^*(t)]$ has the following two properties:*

$$(1) \quad \theta(t = 0) = 0,$$

$$(2) \quad \forall t, t < \theta^{max}, \lim_{t \to 0} \theta(t) = 0.$$

While the scarcity of quick review does not always benefit larger, established firms, reflection suggests that, given a distribution of firms—existing and potential entrants to a market—tardy review will impose heavier *relative* costs on smaller firms. Smaller firms are usually newer ones, and delayed market entry can impede a cash flow crucial to firm survival and cripple the firm's relationship with potential investors who seek certainty over future product revenues. Larger firms are better able to survive this sort of exposure.

THE VALUE OF FAMILIARITY: FIRM EFFECTS AND FREQUENT SUBMISSIONS

The optimal stopping problem in Eq. (5) is repeated numerous times over the lifetime of a regulatory policy. In pharmaceutical regulation, for instance, the FDA has since 1950 received over 10,000 new drug applications (NDAs) and has approved over 1,500 new chemical entities, or drugs with new and complex molecular structures. State public utility commissions face repeated licensing applications for construction and operation of power plants from the same utility companies. For hydropower plants the Federal Energy Regulatory Commission (FERC) conducts such reviews, again repeatedly interacting with many companies. In numerous industries, regulators receive new products or license applications repeatedly from the same firms or interests. It is useful, then, to consider the dynamics of the agency's approval policy when the problem in Eq. (5) is repeated finitely.

If firms differ in the quality and care with which they produce products, then it is reasonable to believe that these differences are at least partially observable by the agency. Anecdotal evidence coheres with this intuition. The FDA employs hundreds of inspectors to survey production processes at pharmaceutical plants, and firms that are found in violation of federal regulations are frequently revisited.[11]

To model the possibility of systematic firm differences in product quality or hazards, I introduce a firm-specific error term to the danger variable. Whereas before the specific danger of a product $i(\mu_i)$ was a random draw from a normal "production distribution" with mean m and variance s, I now respecify μ_i to depend on an independent firm effect, ξ_k, which is distributed standard normal across firms. The danger of any product is now

$$\mu'_{ik} = \mu_i + \xi_k, \quad \mu_i \sim N(m, s), \quad \xi_k \sim N(0, 1).$$

Notice that the parameter m is still the central tendency of the danger distribution, as $E[\mu'_{ik}] = m$. By the addition of independent variances, however, $\mathrm{Var}[\mu'_{ik}] = s + 1$. The rational agency learns about ξ_k, but this cannot be done within the confines of the problem in Eq. (5). It must be done across repetitions of the stopping exercise, that is, *across products*.

Because the danger effect ξ_k is distributed normally across firms, the agency can form a Bayesian estimate $\hat{\xi}_k[\tau]$.[12] Let τ be the "historical time" elapsed since the first firm submitted the first product to the agency, and let $N_k[\tau]$ be the number of products submitted by firm k in historical time, with $M_k[\tau]$ the number of products approved, and $N_k - M_k$ the remainder that have not (yet) been approved. Then the regulator's best estimate of ξ_k is a weighted sum of the danger estimates of any firm's submissions (approved and not approved), as follows.

$$\hat{\xi}_k[\tau] = \left[\left\{ \frac{1}{N_k} \sum_{i \in [N_k \sim M_k]} \hat{\mu}_{it} \right\} + \left\{ \frac{1}{N_k} \sum_{i \in [M_k]} \mu_i \right\} \right] - m.$$

The agency's uncertainty about a firm's safety is the posterior variance $\mathrm{Var}^p[\hat{\xi}_k]$ of the estimated firm effect $\hat{\xi}_k[\tau]$. The prior variance of the firm effect ξ_k is 1, which holds for all firms, but the posterior variance of the estimated firm effect depends on the firm's previous submissions, so that that $\mathrm{Var}^p[\hat{\xi}_k]$ is a decreasing function of N_k.

$$\mathrm{Var}^p[\hat{\xi}_k] = \left[1 + \left\{ \sum_{i=0, k=K}^{N_k - M_k} \frac{r(t_{sub(ik)})}{s} + M_k \right\} \right]^{-1}, \quad \textbf{(7)}$$

where $r(\cdot)$ is the product-specific precision of $\hat{\mu}_t [= s - S(t)]$, and t_{sub} is the review time for all nonapproved products—or the elapsed time from the submission of those products to the present time (τ). When the regulator has had no experience with firm $k (N_k = 0, M_k = 0)$, $\mathrm{Var}^p[\hat{\xi}_k]$ simply reduces to 1, or the "prior" distribution of $\hat{\xi}_k[\tau]$. Given some experience with any firm ($N_k > 1$), then at the beginning of review

[11] For evidence on firm-specific updating in food regulation, see Hinich and Staelin 1980.

[12] I note here that the agency's estimate is not fully Bayesian, because once product review starts, the agency is assumed to learn only about the product under consideration. Hence the agency is not using all of the information available to it as review progresses; knowledge about the firm's other products and expectations about present and future products in the "pipeline" are both fixed at the beginning of review (see also Eq. [8]). This is an absolutely necessary constraint for solution of the model, as relaxation would require that the smooth pasting condition for solution satisfy a highly nonlinear set of constraints. I am currently exploring ways to relax it.

for any given product, the agency has better information about the firm-specific danger effect ξ_K for firm $k = K$. This fact forms the central intuition underlying the advantage of familiarity.

Proposition 5. The Advantage of Familiarity. *For any set of firms $k = 1, 2 \ldots K$ such that $\xi_1 = \xi_2 = \cdots = \xi_K$, the following two statements hold*:

(a) the expected review time $E[t_{app(ik)} \mid \xi_k]$ conditioned on the posterior variance $Var^p[\hat{\xi}_k]$ is a strictly decreasing function of N_k,

(b) the expected review time $E[t_{app(ik)} \mid \xi_k]$ conditioned upon the sample variance of the danger estimate $(\hat{\xi}_k)$ is decreasing in N_k unless $(\mu_{N_k+1} - \bar{\mu}_{N_k})^2 - \hat{\sigma}^2_{\bar{\mu},N_k}/N_k > \Phi(\delta A - m, s)1/N_k^2$, where $\hat{\sigma}^2_{\bar{\mu},N_k}$ is the mean square error of $\bar{\mu}_{N_k}$.

Proposition 5 specifies a relationship between the role of uncertainty about firms and drugs in the approval policy and the role of experience in reducing this uncertainty. If we form general expectations over regulatory behavior, then this uncertainty is understood in terms of the posterior variance of the danger estimate, $\text{Var}^p[\hat{\xi}_k]$, and the relationship between expected approval times and regulatory learning is monotonic. If uncertainty is understood as the sample variance of the danger estimate, then for most realizations of μ_{ik}, expected approval times will decrease with regulatory learning.[13] Extreme realizations of μ_{ik} may increase regulatory uncertainty. However, the left-hand side of the inequality in Proposition 5(b) has zero expectation, so these jumps in regulatory uncertainty will not happen often. And they are less likely to happen when N_k is small.

Because the regulator's uncertainty about a firm is a generally decreasing function of regulatory experience with that firm (N_k), so is uncertainty about the drug under review, or $S_{i(k)}(t)$. Since the optimal policy requires product approval when estimated danger $\hat{\mu}_t$ hits the barrier $\eta*(t)$ *from above*, an increase in familiarity (N_k) pushes $S(t)$ lower—and thus the barrier higher—for all $t > 0$. Because the optimal stopping policy established in Proposition 1 invokes risk-neutrality, this is a strong result. Reduced *prereview* uncertainty about μ_{ik} translates into quicker approval times, conditioning on the event of approval.[14]

Notice that the result in Proposition 5 makes no assumptions about how attributes of firms might affect the approval payoff. The only differences that prevail among firms are captured in N_k, the number of submissions. Over any given period of time, larger firms will submit more products than smaller ones, and *this difference alone is sufficient to yield lower expected approval times*. Hence there is protection—smaller and newer firms will, in expectation, wait longer to enter any given market—without capture.

The second result is more powerful. Because the result is conditioned on the firm-specific danger value $\hat{\xi}_k$, Proposition 5 raises the possibility that firms with many submissions and with positive $\hat{\xi}_k$ (or *below*-average safety reputation) *may actually receive quicker reviews*, conditioning on the event that a given product is approved. In other words the advantage of familiarity outweighs poor firm reputation, to a point. I formalize this in Corollary 5.1.

Corollary 5.1. *Consider two products with identical experimental histories and identical approval payoffs but submitted by different firms: one $(k = k^0)$ with no previous submissions, the other $(k = k^B)$ with a "bad reputation" $(\hat{\xi}_{k^B} > 0)$ but positive familiarity $(N_{k^B} > 0)$. Let t_{stop} be the approval time for the drug submitted by the unknown firm (k^0). Then when*

$$\hat{\xi}_{k^B} \leq \frac{\sigma_i^2 + st_{stop}}{\sigma_i^2}[S_{k^0}(t_{stop})^2 - S_{k^B}(t_{stop})^2]$$

or (conceptually) when drugs receive sufficiently "early approval" or the bad firm has a sufficiently low firm danger estimate low $\hat{\xi}_{k^B}$), the product submitted by the bad but familiar firm receives quicker approval than the product submitted by the unknown firm.

One example of a perfect (one-for-one) tradeoff between the disadvantage of a bad reputation and the advantage of familiarity appears in Figure 2. Here a bad reputation—higher values of $\hat{\xi}_{k^B}$—boost the starting point of the process $\hat{\mu}_t$ from m to $m + \hat{\xi}_{k^B}$ and slow approval. Yet to a limited extent, familiarity advantages can compensate and even outweigh these effects. Even as a bad reputation moves the process $\hat{\mu}_t$ away from the barrier, familiarity shifts the barrier back toward the process.[15]

EARLY–ENTRANT PROTECTION

Another characteristic of large firms is greater capitalization and an ability to enter new market niches at lower cost. The first new entities for AIDS, acute coronary syndrome, and erectile dysfunction were all developed by large companies (e.g., Burroughs–Wellcome

[13] I thank an anonymous reviewer for pointing out this alternative notion of regulatory uncertainty.

[14] An important caveat merits remark. As Dixit (1993, 58) warns, it is *not* always the case that reduced uncertainty about a Wiener process translates into a quicker optimal stopping time. Recall that a crucial determinant of the agency's uncertainty about μ_{ik} is σ. As Dixit adroitly notes, increases in σ may actually *reduce* the stopping time because σ represents the volatility of $X(t)$ and a more volatile process may more quickly reach the barrier. Proposition 4 rests on a different result and is robust to Dixit's cautionary note. Recall that *σ does not differ across products but that the posterior variance of ε_k does*. At the beginning of any product review $(t = 0)$, the agency knows the "prior" variance of μ by $S(0) = \{s + \text{Var}^p[\xi_k]\}^{-1}$. In finite time, $S(t)$ is *always* an increasing function of $\text{Var}^p[\xi_k]$, independent of the value of σ.

[15] If product submission were costless, then rational firms might attempt to take advantage of this familiarity result and load the regulator with frivolous submissions to increase their familiarity with the regulator. As DiMasi et al. (1991) show, however, pharmaceutical product development is enormously expensive even before the regulatory approval stage, so even if firm submissions were endogenous to the model, it is doubtful that this would be an equilibrium policy under any reasonable submission cost parameters.

FIGURE 2. Trade Off between Firm Danger Reputation and Regulatory Familiarity

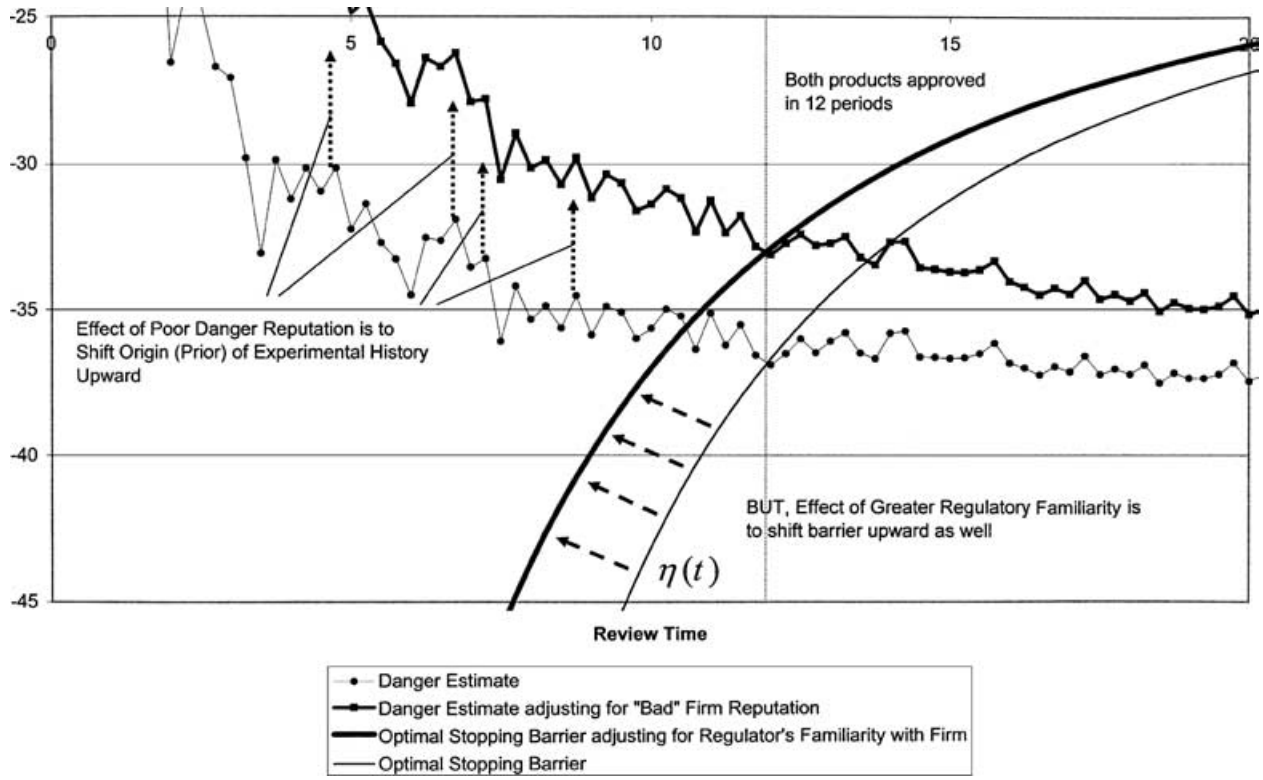

with AZT). DiMasi et al. (1995, 213) present evidence consistent with the proposition that larger firms enter new markets earlier, as their drugs are more likely to receive FDA priority status and their per-drug sales are higher than those of small firms (see also Dranove and Meltzer 1994, 409). If earlier drugs for a disease receive quicker approvals, then larger and older firms benefit disproportionately if in fact they enter these new markets sooner than smaller firms do.

In this section I adjust the agency's approval function to consider the effects of the political demand of consumers on firm advantages. I focus on diseases as "market niches" because they play a crucial role in the pharmaceutical industry. In essence, there is no single "market" for pharmaceuticals; there are, rather, numerous markets bounded in size by the diseased population. This section of the argument is therefore tailored much more specifically to the case of drug regulation. (I consider other possible applications below.)

In the past two decades, the FDA has shown considerable responsiveness to pressure from organized patients and disease-specific advocacy groups.[16] The most visible case of FDA reaction to these forces was the quickening of drug approvals for AIDS drugs in the 1980s and early 1990s (Epstein 1996). More generally, as Olson (1995, 404) argues, "From the FDA's perspective, the threat of adverse feedback from consumers seems to dominate the complaints of the drug industry." In this section, I specify the agency's approval payoff as a function of the political organizations of consumers (and of producers). I also provide a functional form for the effect of past and future drugs on the approval payoff for the present drug. The approval payoff may be defined intuitively as follows.

> For any drug i, which treats disease j and is submitted by firm k, the agency's payoff from approving the drug (denoted A) is equivalent to the sum of all individuals with disease j who have no available pharmaceutical alternatives and whom drug i would be expected to cure, where each consumer is weighted by their relative political organization (a political multiplier), and where the drug itself is weighted by the firm's political clout.

Before translating this principle into an equation for the approval payoff, I pause to note several features of this concept. I assume that the political demand for a drug differs from the economic demand for it in an important way. Political demand is greater for those individuals who have *no* therapeutic alternatives for their disease. In other words, if individuals *are taking* a drug that ameliorates their condition in some way, then they have *less* political demand for any more drugs for their disease, even if these drugs would improve their condition or would be available at a lower cost. In other words, the model rests on the assumption that, once

[16] Some of these are connected to the pharmaceutical industry; the vast majority are not.

individuals have adopted a drug, their contribution to the political demand pool drops to zero.[17]

In short, while individuals as consumers may maximize their health status and care about price, *individuals as citizens satisfice with respect to the agency*. The reason, I suggest, is that citizens have "traceability" constraints (Arnold 1990). They do not blame the agency for the high price of drugs—even if it is apparent that the regulatory process boosts drug prices. If the price of a drug is too high, or if citizens are weakly satisfied with a drug they are taking, they blame not the agency, but the producers.[18]

The Effect of Past and Future Drugs

Because political demand is a function primarily of those patients *without* therapeutic alternatives, the payoff is not a simple function of the curing power γ_i of the drug submitted. The reason is that drugs may already exist to treat disease j, and forthcoming drugs in the "pipeline" may also be expected to treat the disease. Existing drugs influence the approval payoff in the following way. If no drug has previously been approved for disease j, the total expected curing of drug i is $\gamma_{ij} L_j$. If one drug ("drug 1") has already been approved, then the expected curing of drug 2 would be $\gamma_{2j} L_j \times (1 - \gamma_{1j})$. So for any series of drugs $i = 1, 2, \ldots, N$, the total curing of the Nth drug, given past approvals, is

$$\gamma_{Nj} L_j \times (1 - \gamma_{1j}) \times (1 - \gamma_{2j}) \times \cdots \times (1 - \gamma_{N-1,j}) = L_j \{\gamma_{Nj}\} \prod_{i=0}^{N-1,j} (1 - \gamma_{ij}).$$

Of course, pharmaceutical firms constantly introduce new drugs, and the rational agency will take into account this stream of future submissions. If the agency expects a sufficient number of high-quality drugs to be submitted for disease j in the very near future, then the approval payoff for the current drug is lower. To incorporate into the approval payoff the set of drugs that *will* be submitted and *may* be approved in the future, define χ as the time that the Nth drug (the one currently under consideration) is submitted. Let E_χ be the expectation operator computed at time χ. Let c_{ij} be the mean of the curing distribution $f(\gamma_{ij})$. Then the payoff for the Nth drug (A_N) is

$$A_N = L_j \psi_j \rho_k \left\{ \gamma_{N_j} \prod_{i=0}^{N_j - 1} (1 - \gamma_{ij}) \right\} \times \left\{ \prod_{i=N_j+1}^{N_{\max}(\lambda)} \left[1 - e^{-\delta(E_\chi[t_{sub(i)}(\lambda)] + E_\chi[t^*_{app(i)}])} c_{ij} G^*_\chi(A_i, c_{ij})\right] \right\}. \quad \textbf{(8)}$$

The second term in brackets on the right-hand side is called the *pipeline value* and will be denoted A_π. The pipeline value is the expected fraction of present prevalence that remains after the curing of the full stream of future drugs, weighted by the probability of their approval under the optimal policy, and each fully discounted. The pipeline value can be zero only if the agency *expects* a cure-all to be approved with probability 1 immediately upon the submission of the drug. By Proposition 2, this event has negligible probability.

I also allow the approval payoff to be a function of the political influence of the firm submitting the drug, parameterized as ρ_k. (The first four parameters in Eq. (8) are assumed positive.) This parameter embeds the idea that better organized interests will receive preferential treatment. Capture theory may not always imply such a relationship, but Stigler's opening epigraph certainly does.

Intuitively, "early-entrant protection" (EEP) is the systematic advantage in expected time to approval for the first drugs approved for a disease relative to drugs developed later. The premise that drug consumption "removes" patients from politics is sufficient to demonstrate two forms of "protection" for the manufacturers that produce the first drugs for any disease.

Strong early-entrant protection: Expected approval times are a strictly increasing function of the order of entry, such that $E[t_{app(i=N+1)}] > E[t_{app(N)}]$ for all N.

Weak early-entrant protection: One drug has a lower expected approval time than all subsequent drugs.

Analysis of the model suggests that early entrant protection depends critically upon how the curing power of therapies unfolds over a sequence of drugs targeting a given disease. An instructive result comes when we assume that this unfolding is stationary, that is, it has constant expectation or mean.

Proposition 6. Concave Early-Entrant Protection Under a Stationary Quality Distribution. *For any disease $j = J$, consider a sequence of disease-specific drugs $i = 1, \ldots N_j$. Then $\gamma_{i,j}$ is the curing probability for drug i for disease j. Let $f(\gamma)$ be a stationary curing distribution for any disease, such that $E[\gamma_i] = g$ for all $i \in J$. Then strong early entrant protection holds, such that*

[17] One can relax this assumption by assuming that the political demand contribution drops by a factor of ζ, $0 < \zeta \leq 1$. The general results here are unaffected by this transformation. I opt for the simpler variant here. While this assumption may seem constraining, it characterizes the politics of pharmaceutical regulation much better than any alternative. There are few documented cases of citizens or their representatives lobbying the FDA for quick approval of drugs that are cheaper for a given disease than existing drugs. More significant, there are *many more* cases in which new, better, and cheaper drugs are submitted to the FDA and in which no lobbying for approval by citizens occurs. Also, I am unaware of any instance in which the FDA has cited the price-competition value of a new drug in approving or rejecting it.

[18] The story of AIDS drugs is instructive and supports this assumption. AIDS protesters swamped FDA headquarters when the drug ddI was under review (Hilts 2003). Accordingly, as the model here predicts, the agency approved the drug quickly. Yet when AIDS patients were distraught at the high initial price of AZT and other AIDS drugs, they protested not at the FDA, but on Wall Street, where ACT-UP members chained themselves to the floor of the New York Stock Exchange. Burroughs–Wellcome subsequently cut the price of AZT from $8,000 to $6,400 (Epstein 1996).

FIGURE 3. Early Entrant Protection as a Function of Firm Political Clout

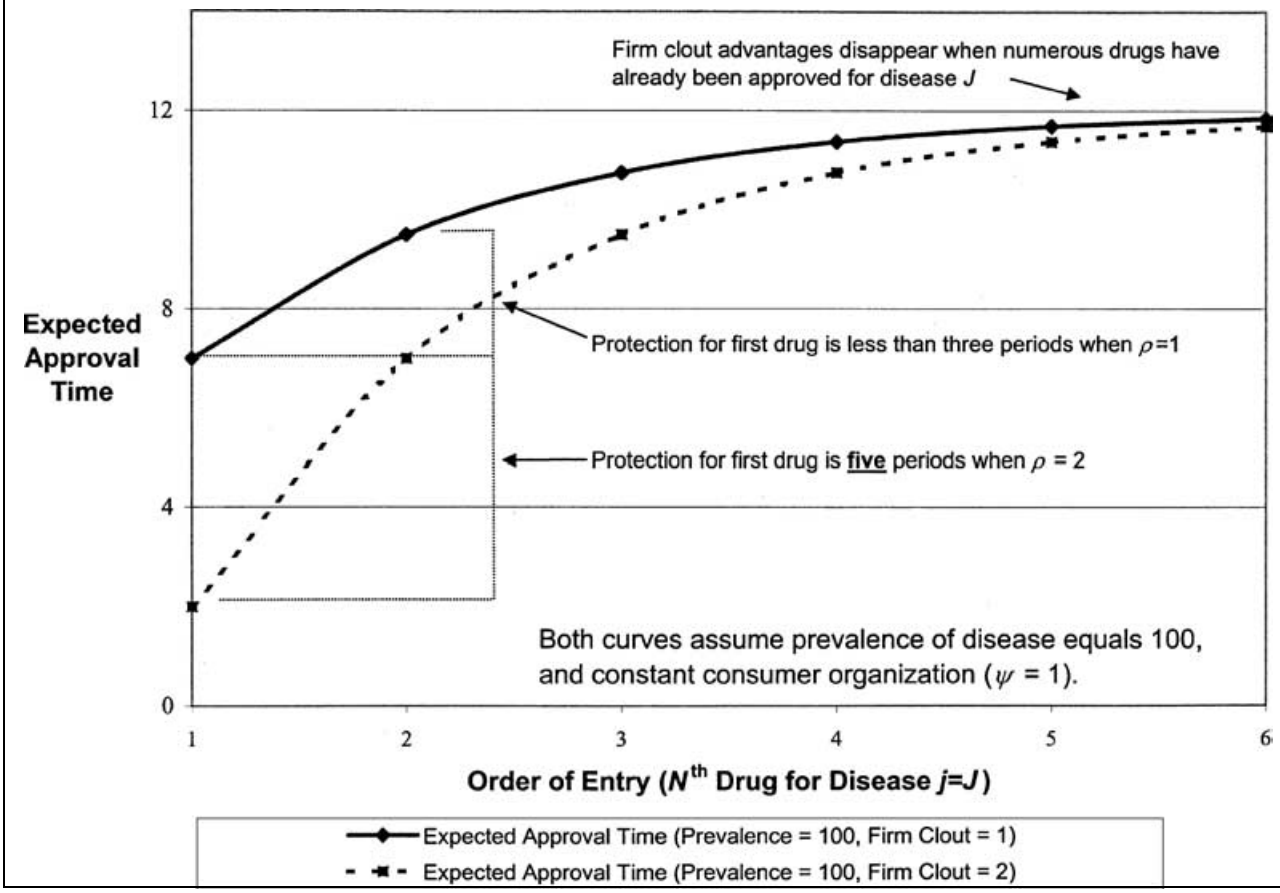

$E[t_{app(N+1)}] > E[t_{app(N)}]$ *for all i. Moreover, the degree of protection — the approval time advantage of earlier as opposed to later entrants—is concave in* N_j.

Concavity is far from trivial. It means not only that the first entrant to any disease niche receives the quickest review, but also that *the second entrant has the largest expected delay relative to the previous entrant*. In other words, the degree of protection given to the first entrant is the greatest.

Possible Interactions between Capture and Noncapture Factors: The Effect of Firm Clout. The case of stationarity is instructive for several reasons. The degree of early-entrant protection that a firm receives is an increasing function of *all* of the other parameters of the approval payoff. In particular, firm clout (ρ_k) boosts the effect of early entrant protection, and vice versa. This result is depicted in Figure 3, where the concavity of the protection function also appears. Where large firms have political rents and can influence the agency, it is a genuine advantage for such firms to enter new markets early, as their competitors will have disadvantages both in influence and in reduced "political demand" for new drugs.

The Generality of Early-Entrant Protection. More generally, for any two drugs $i = N$ and $N + 1$ with identical levels of danger ($\mu_N = \mu_{N+1}$), the expected approval time for the Nth drug is always shorter unless the $(N + 1)$th offers an improvement in curing power.

Proposition 7. *Given two drugs with identical danger* ($\mu_N = \mu_{N+1}$), *then unless the* $(N + 1)$*th drug improves upon the curing of the Nth, EEP must hold. The inverse is not true.*

Early-Entrant Protection When Drugs Improve Over Time. To assume the stationarity of $f(\gamma_{ij})$ is restrictive; it implies that the curing power of drugs does not improve as new ones are developed. It is useful, then, to consider the case of *stochastic improvement*, where newer drugs will be expected to have superior curative power relative to older ones, a process driven by technological change.

Once we admit stochastic improvement in $f(\gamma_{ij})$, then early-entrant protection does not hold for all sequences of products. Intuitively, if the regulator expects the first drug for any disease to cure 10% of patients and firmly expects the second drug to cure 90%, the agency will consistently find it worthwhile to approve the second more quickly. A generalizable result, then, requires limits on the path of γ_{ij}. A realistic assumption is *concave* improvement, which would reflect decreasing marginal returns to technology or R&D.

Proposition 8. Early-Entrant Protection for Concave Improvement Functions Over the Curing Distribution. *Let the stochastic improvement function satisfy (1)* $d\mathrm{E}[\gamma]/dN > 0$ *and (2)* $d^2E[\gamma]/dN^2 < 0$, *for all N. Then there are at least two forms of early-entrant protection.*

(1) For γ_1 *sufficiently high, there is strong early-entrant protection.*
(2) For any $\gamma_1 < 1$, *and with maximal improvement consistent with concavity, there is always weak early-entrant protection, or protection for some set of drugs submitted and approved early in the sequence.*

Assuming decreasing marginal returns to technology, simple restrictions on the improvement of the second drug over the first (or the third drug over the second) are sufficient to derive shorter expected review times for the first drug or drugs approved, *relative to all other drugs*. Proposition 8 shows that early-entrant protection holds even when the regulator *expects* later entrants to a market niche to offer quality improvements.

Corollaries: First Drugs, Innovation, and Product Diversity

A number of other testable statements about early-entrant protection (EEP) can be advanced here.

Corollary 8.1. *The greater the quality of the first product approved, the greater the likelihood of EEP for product* 1 *relative to the sequence of following products.*

Corollary 8.2. *For any sequence of products, the likelihood and severity of EEP are decreasing in the product innovation rate* λ_j.

Corollary 8.3. *For any firm, the likelihood of receiving EEP for at least one product is increasing in the diversity of market niches targeted.*

Corollary 8.1 follows rather straightforwardly from the specification of the approval payoff in Eq. (8). If, for instance, the first drug to be produced and approved for a given disease cures 90% of the prevalent sufferers, the agency will not approve any drug as quickly given the drastic reduction in *marginal* curing for later drugs.

The Effect of Product Innovation. Corollary 8.2 suggests that the potentially distortionary effects of early-entrant protection will be mitigated when the submission rate is high. The reason is that, when the first drug (or first few drugs) for any disease is being considered, the agency will discount their marginal curing value by the expectation that other cures are likely to be submitted soon. One policy implication is that policy instruments that enhance the climate for product innovation may, ceteris paribus, reduce the advantage of early entrants in the regulatory process.

The Corollary Advantage of Diversity. If larger firms have greater product diversification (products in more markets), then Corollary 8.3 points to another advantage for them. All else equal, firms with drugs for more markets have a greater likelihood of receiving protection in an least one market niche from the regulator. Smaller drug firms tend to concentrate on one or a few diseases. Unless they enter these niches quickly, their relative regulatory disadvantages will be severe.

Other Approval Processes Where Early-Entrant Protection May Apply. While I have tailored the discussion of early-entrant protection to pharmaceutical markets, it is quite possible that the model here (with elaboration or revision) could apply in other situations. Consider the licensing of alcohol service in a small town where there is consumer demand for just two types of establishments: an Italian restaurant (wine and grappa) and a microbrewery (beer). Suppose that local political demand for the types is equal and that, over a relevant period of time, all applications for an alcohol license offer products with identical expected quality (revenues) and identical expected hazards (loud and drunken patrons). If the township licensing regulator has discretion over licensing decisions, early-entrant protection dynamics suggest that the first microbrewery to apply for an alcohol license will receive more favorable treatment than either (a) the tenth Italian restaurant to apply or (b) the second microbrewery to apply (once the first microbrewery has been approved). The reason, in the logic of early-entrant protection, is that local *political* demand for additional alcohol licenses will have been nearly satiated by the appearance of a few Italian restaurants and one microbrewery.

Space does not permit a careful analysis of other applications, but allow me to trace (very speculatively) the following possibilities: (1) Might a transportation regulator approve the first proposed flight route from Chicago to Peoria more likely and more quickly than (a) the tenth route from Chicago to Washington or (b) the *second* route from Chicago to Peoria? (2) Suppose that in a three-year period a political science journal has published 10 political theory articles but not a single article on international relations. Might a journal editor (who is facing increasing criticism from his or her colleagues in security studies) approve the next decent IR submission with greater likelihood and rapidity than (a) the equally good "eleventh" submission in political theory or (b) the third or fourth new IR submission, once the first IR submission in three years has been published?

THE PROBLEM OF ENDOGENOUS SUBMISSIONS

The principal constraint of this model is its assumption that firm submissions are exogenous. Endogenous submissions may pose two potential problems. First, the regulator may face trivial uncertainty if the act of regulatory submission is itself a signal that completely reveals product quality or hazards. If the regulator's standards were high enough and firms submitted only

"good" products submitted in equilibrium, then there would be little point in extended regulatory review. So one question is: Would regulatory uncertainty disappear or substantially wane in a strategic setting? Second, even if considerable uncertainty remained, the regulator would undoubtedly know that some sorts of submissions were unlikely, and hence the distribution of product hazards for observed submissions would likely be truncated or characterized by an upper bound. Do the results here hold for product hazard distributions of this character?

The Persistence of Regulatory Uncertainty

As long as firms retain some private information about the quality or hazards associated with their product, significant regulatory uncertainty would persist in a strategic setting with endogenous product submissions. A fuller analysis of this question lies beyond the present paper (see Carpenter and Ting 2004). The essential intuitive point is that, in any game with endogenous submissions, there exists no semiseparating equilibrium in which only types above some threshold submit and the regulator approves all submissions; the regulatory process causes a partial "pooling" of products. (That is, there are not equilibria where low-quality products are always abandoned, and only good products are submitted, and in which the regulator uses this information to relax the regulatory process.)[19] The reason is that, if by dint of high regulatory standards only good products were submitted, then the regulatory agency would find it worthwhile to shortcut review and approve all products immediately. But this move would induce firms to submit "bad" products, and so the regulator must randomize review and approval to a degree. This means that a substantial portion of the applicant pool will in fact be below the regulator's standards, even though the regulator can never determine with absolute certainty which products these are. So while the act of submission provides some information to the regulator, considerable uncertainty remains.[20]

Extension to "Reduced Form" Truncated Prior Distributions

Do the results hold when prior distribution is truncated (i.e., certain submissions are not possible)? Briefly, the answer is yes. Consider first the advantage of regulatory familiarity. Recalling that the prior distribution of $\hat{\mu}_t$ is normal, we can represent the impossibility of certain submissions by an upper-truncated normal distribution. The question is then whether the agency can update on such a distribution in the way it does with the normal distribution in Proposition 5. As with the normal distribution, the posterior variance of a *truncated* normal distribution is also decreasing in sample size, hence a result much like Proposition 5 would extend to the case of a truncated prior as well. Second, note that none of the early-entrant protection results depend on the distribution of $\hat{\mu}_t$ or $\hat{\xi}_k$, so these results extend.

Theoretical and Empirical Reasons for Decision-Theoretic Specification

Finally, a decision-theoretic analysis of regulatory review has clear analytic value on its own. First, considerable contextual and informational complexity can be added to the model, richness that must be sacrificed in a tractable analytic game. Specifically, a tractable game-theoretic model of approval regulation admits of only two product types, allows no differentiation between quality and hazards, permits no regulatory review of evidence, and allows for no temporal discounting or dynamics, no role for organized consumers, no repeated interaction between firms and regulators, and no possible entry-order effects (Carpenter and Ting 2004). It is for this reason that a decision-theoretic model, with its own limitations, is uniquely useful in understanding many dynamics that a game-theoretic model with endogenous submissions must sacrifice. A second reason is empirical; regulatory review happens 100% of the time a product is submitted, and observers note relatively high rejection rates of drugs submitted (this is impossible to know with certainty, but the GAO [1995] estimates that 40% to 50% of new drug applications are rejected at the FDA). We can also document facts consistent with considerable uncertainty in agencies that make such decisions (reviewers wanting more time, and asking for more information, and split votes on advisory committees). While game-theoretic analysis would complement this paper, then, it remains true that the *combination* of game- and decision-theoretic analysis (at least at present) exceeds the value of game-theoretic analysis alone.

BOUNDED RATIONALITY: ADAPTATION, PLACEBOS, AND FORGETTING BY TURNOVER

The model developed here assumes a good deal of rationality; the agency behaves as if it were solving a complex dynamic programming problem, it conducts real-time Bayesian updating on a continuous stochastic process, and it knows and uses both past regulatory submissions and (expected properties of) future regulatory submissions. It is worth reflecting, then, on whether the results derived here would hold under the more realistic scenario of a boundedly rational regulator. While each of the following ideas would require separate papers to elaborate, there are at least three ways that such bounded rationality could be approached.

[19] I thank an anonymous reviewer for this language.

[20] In the semiseparating equilibrium of Carpenter and Ting (2004), the firm mixes between submitting and abandoning. The regulator sets approval so that the "low" type is made indifferent between submitting and withdrawing. A similar semiseparating equilibrium occurs in many crisis bargaining games studied in international relations (Schultz 1998).

A first possibility is that the real-world agency's learning is not perfectly Bayesian but rather "adaptive" (e.g., Bendor, Mookherjee, and Ray 2001). Instead of using Bayes' rule to update firm histories, for instance, the regulator might keep a running tally of experience with a firm in which only sufficiently "bad" interactions enter into memory. Or the regulator might be subject to "confirmatory bias" (Rabin and Schrag 1999), whereby the first interaction with the firm strongly conditions inference from subsequent interactions. The key to familiarity advantages is that, as long as repeated experience increases (even slightly) the precision with which the regulator can estimate firm danger, the familiarity result can still obtain. Hence the familiarity result probably extends to cases of non-Bayesian updating but may not hold under confirmatory bias.

Second, patients lobbying the agency may be subject to "placebo learning," that is, overestimation of the quality of a product by virtue of placebo effects. FDA approval is sometimes influenced by the emotive appeals of patients who testify that the drug worked for them and that they had no other option. Even if the drug is effective, some of these patients' responses may be placebo-related. If the agency incorporates such lobbying without discounting for this possibility, it may overestimate the product quality (curing power) and underestimate product hazards. In other words, bounded rationality on the part of organized interests *may* translate to bounded rationality on the part of the regulator.

Third, personnel turnover may create a regulator with imperfect memory. The model here assumes a regulator with continuous and perfect memory; the regulator remembers not only the entire history of the present product review, but also the history of each firm that submits. In real-world settings, memory limitations strongly influence optimization decisions (Mullainathan 2002). In reality, new bureaucratic officials arrive to agencies without built-in memories (the histories and cases must be learned), and officials who depart the agency take valuable firm-specific and case-specific information with them (which may not be immediately or costlessly transferable to the officials who replace them). Given the pervasiveness of turnover in government workforces, turnover as "organizational memory" becomes a crucial policy issue. As a conjecture, it is quite possible that large firms whose reputations are well known to the general public and scientific community may suffer less when turnover is high, because their external reputations "compensate" or "substitute" at some level for the regulator's loss of information. Hence organizational memory might in fact exacerbate large-firm and older-firm advantage in regulatory settings.

CONCLUSION AND EMPIRICAL IMPLICATIONS

Conventional wisdom follows Stigler in arguing that, upon observing distinct policy or regulatory advantages for apparently wealthy or well-organized firms, one should infer that capture, rent-seeking, or some other form of distributive politics is necessarily at play. Yet Stigler was wrong: evidence of policy advantage is quite insufficient for a conclusion supportive of capture theory. A neutral learning regulator motivated only by reputation protection and constrained by political responsiveness to consumers would also provide advantages to larger and older firms.

While it generates some results that are observationally equivalent to those of capture theory, the reputation and learning model also has some divergent empirical implications. First, Proposition 5 and Corollary 5.1 suggest that observed large-firm advantage in approval regulation should attenuate (perhaps entirely) once previous submissions are controlled for. Evidence for a reduced apparent firm size effect, controlling for familiarity, would be evidence against the capture account and evidence in favor of the model here.[21] Second, the model predicts that firm reputations for safety will decisively influence regulatory outcomes; capture theory does not offer such a prediction. Third, familiarity effects should be conditioned both on the length of the review (see Corollary 5.1) and on features of the regulatory bureaucracy itself (the institutional memory of the organization, for example).

While capture theory might also predict early-entrant advantage in approval regulation, there are numerous predictions made by the model for which capture theory cannot account. Early entrant protection should be conditioned not on the political organization of producers, as capture theory might predict, but on the political organization of *consumers*. Hence the model here uniquely predicts interaction effects between order of entry and consumer political organization. In addition, the corollaries of Proposition 8—–the importance of the quality of the first product on the market and the regulators' adaptation toward the pipeline of future products—would also be inconsistent with capture theory.

Recent studies have arrived at empirical results supportive of the reputation and learning perspective advanced here. Olson (1997) interprets firm characteristics as potential signals to uncertain regulators about unobserved product quality. The formalization here generally supports Olson's argument and, conversely, Olson's results provide empirical support for this model. Kyle's (2002) results also provide support for early-entrant protection and familiarity advantages. The regulatory advantage of familiarity may best be interpreted not as political "coziness" between regulator and industry (Bartel and Thomas 1987; Stigler 1975, 162–63), but as the result of repeated observation of a sequence of firm products by a rational agency.

Familiarity advantages can hold even when an uncertain regulator encounters firms with a below-average safety record. While not all such firms have regulatory

[21] The proper framework for test would be to assess the coefficient on firm size in a duration model of product approval times, and to examine whether the estimated marginal effects of firm size decline in the presence of controls for regulatory familiarity. Work is under way in this direction. I thank an anonymous reviewer for this suggestion.

advantages, some of them will still receive quicker approvals when their products are accepted. As a result, the model here departs in some respects from Olson's logic; familiarity advantages are distinct from (though not independent of) firm-specific quality differentials. A test distinguishing between Olson's predictions and mine would require data on the actual quality of a product. The model here predicts that, even controlling for product quality, a familiarity effect should hold such that a firm's expected approval times should be decreasing in the number of previous product submissions.

The reputation and learning model may also explain patterns of large-firm advantage that prevail in areas of regulation outside of the pharmaceutical industry. These patterns include quicker license approvals for older and more established firms in the electricity industry and quicker permit granting for and quicker disposition of litigation for older and larger firms in the enforcement of environmental regulation (Gordon 1999; Spence 1999).

In conclusion, the reputation and learning model suggests that even an apparently "neutral" regulatory regime can bequeath systematic advantages to larger, more established firms. This result should give pause to students of institutional design who believe that rigorous political insulation of an agency can create an even playing field for all firms large and small. In this light, the model is similar to other formal efforts demonstrating the impossibility of neutral administrative arrangements (Hammond and Thomas 1989).

The lesson for students of political economy is one of proper inference. Observed advantages for larger and older firms in regulation may say little or nothing about politics because they may exist for many reasons, some related to capture and some not. Whether or not capture prevails in a given industry, the advantage of familiarity and early-entrant protection may interact with capture to magnify large-firm advantage in regulated industries.

APPENDIX: PROOFS

Proof of Proposition 1[22]

The existence of an optimum for (5) has been shown elsewhere (see, for example, Shepp 1969 and Miroschnichenko 1975). Here I show that the optimal barrier is Eq. (7). The solution to Eq. (5) must satisfy the functional (Bellman) equation,

$$\delta F(\hat{\mu}, t) = \{E_{\hat{\mu},t} F(\hat{\mu}(t+dt)) - F(\hat{\mu}(t))\} + o(dt), \qquad \textbf{(A1)}$$

where $o(dt)$ represents vanishing terms of order higher than t. Dropping higher-order terms and using Ito's lemma,

$$\begin{aligned} \delta F(\hat{\mu}, t) &= \{E_{\hat{\mu},t}[F(\hat{\mu}(t+dt)) - F(\hat{\mu})]\} \\ &= E(dF) = F_t(\hat{\mu}, t) + \frac{1}{2}\frac{S(t)^2}{\sigma^2} F_{\hat{\mu}\hat{\mu}}(\hat{\mu}, t) \qquad \textbf{(A2)} \\ &\Rightarrow F_t(\hat{\mu}, t) + \frac{S(t)^2}{2\sigma^2} F_{\hat{\mu}\hat{\mu}}(\hat{\mu}, t) - \delta F(\hat{\mu}, t) = 0. \end{aligned}$$

[22] A longer version of Proposition 1 appears in Carpenter 2002; I repeat it here for clarity and ease of reference.

Let $[\eta(t), t]$ be the boundary of the continuation region, where the following conditions hold.

$$\begin{aligned} F(\hat{\mu}, t) &= A - (\mu^*/\delta) && \text{(value matching)}, \\ \frac{d}{dt} F(\hat{\mu}, t) &= F_t(\hat{\mu}, t) = \frac{d}{dt}[A - (\mu^*/\delta)] && \text{(smooth pasting)}. \end{aligned}$$

The smooth pasting condition requires identity of the time derivatives for the value function ($F_t(\hat{\mu}, t)$) and the stopping payoff, hence $d[A - (\mu^*/\delta)]/dt = 0 \Rightarrow F_t(\hat{\mu}, t) = 0$. Applying these conditions to Eq. (A2) gives

$$\frac{S(t)^2}{2\sigma^2} F_{\hat{\mu}\hat{\mu}}(\hat{\mu}, t) - \delta[A - (\mu/\delta)] = 0. \qquad \textbf{(A3)}$$

Evaluating (A3) at $\hat{\mu} = \eta(t)$ yields the optimal rule. □

Proof of Proposition 2

For simplicity, fix the barrier η at δA (a process that exceeds δA infinitely often will also exceed every point below δA infinitely often). Then the probability of the process dropping below η is a function of the cumulative normal distribution:

$$\Pr[\hat{\mu}_t^* \leq \eta \,|\, \mu(0) = m] = \Phi\left\{\frac{(\eta - m) - (\mu - \delta A)t}{\sigma\sqrt{t}}\right\}.$$

I now characterize $\lim_{t\to\infty} \Phi_t$. As long as $\mu < \delta A$,[23] the argument of the cumulative normal integral tends to infinity [the numerator increases at a rate of order t; the denominator, at a rate of order $sqrt\ (t)$] and the probability that the process remains above the barrier asymptotically [$=1 - \Phi_\infty(\cdot)$] is zero. If $\mu > \delta A$, however, the argument in $\Phi_\infty(\cdot)$ runs to negative infinity. Then the asymptotic probability of the process resting below the barrier is zero, and the probability of infinite review is strictly positive. □

Proof of Proposition 3

By Proposition 2, analysis must be restricted to the case $\mu < \delta A$, as the case $\mu > \delta A$ yields infinite expectation. Given this assumption and the scale invariance of Brownian diffusions, the first-passage time can be represented as an inverse Gaussian form (Folks and Chikkara 1979), with moment generating function

$$\Lambda_{G^*(t)}[\zeta] = \exp\left[\nu(\delta A - \mu)\left\{1 - \sqrt{1 - \frac{2\zeta(\delta A - \mu)^{-2}}{\nu}}\right\}\right] \qquad \textbf{(A4)}$$

where ν is the scale parameter of the first-passage time distribution (the variance is positive and finite for $\mu < \delta A$), and index variable ζ. The expected approval time can be calculated from the first moment as $(\delta A - \mu)^{-1}$. □

Proof of Proposition 4

Result (1) obtains by the property of $g^*(t)$ such that $g^*(0) = 0$. Result (2) follows from the unique mode of hazard at $\theta_G^{\max} > 0$ (Folks and Chikkhara 1979).

[23] The case where δA equals μ is more complicated. Here the probability of eventual approval is one, but counterintuitively, the expected approval time is infinite due to the positive probability of infinite sojourns of $\hat{\mu}_t$ infinitesimally close to the barrier. See Dixit 1993 for more discussion.

Proposition 5: The Advantage of Familiarity

(a) *$E[t_{app} \mid \xi_k]$ conditioned on the posterior variance* $\mathrm{Var}^p[\hat{\xi}_k]$ *is a strictly decreasing function of N_k.*

Proof. To make the agency's policy feasible, all estimates on the posterior distribution of ξ_k are fixed at the beginning of the review for any drug, otherwise parameters of the decision process (in particular, s) would vary over the course of the review. This assumption makes use of the following lemma.

Lemma 1. *Any $t_{sub(ik)}$ is a optional stopping time for the process $\hat{\mu}_t$, and by the right-continuity of the filtration $\Im_t$, any $t_{sub(ik)}$ is progressively measurable with respect to the σ-field $\Im_{t_{sub(ik)}}$,* as is the stopped process $\hat{\mu}_{t_{sub(ik)}}$.

Proof. See Karatzas and Shreve 1991, Propositions 2.3 and 2.18.

Lemma 1 establishes the feasibility of the agency "stopping" at any point in time and retrieving all current estimates of $\hat{\mu}_t$for all drugs, whether approved or not.

The posterior variance of the estimated firm effect is, from Eq. (7),

$$\mathrm{Var}^p[\hat{\xi}_k] = \left[1 + \left\{\sum_{i=0,k=K}^{N_k - M_k} \frac{r(t_{sub(ik)})}{s} + M_k\right\}\right]^{-1}.$$

This formula presupposes the progressive measurability of $t_{sub(ik)}$. By inspection, $\mathrm{Var}^p[\hat{\xi}_k]$ is strictly decreasing in N_k. Then the agency's submission-adjusted uncertainty over the danger of drug i is $S(t_i \mid \xi_k) = (\{1/[s + \mathrm{Var}^p(\xi_k)]\} + (t_i/\sigma_i^2))^{-1}$, which is strictly increasing in $\mathrm{Var}^p[\hat{\xi}_k]$. ■

(b) *$E[t_{app(ik)} \mid \xi_k]$ conditioned on the sample variance of the danger estimate $(\hat{\xi}_k)$ is decreasing in N_k unless $\{[(\mu_{N_k+1} - \bar{\mu}_{N_k})^2 - \hat{\sigma}^2_{\bar{\mu},N_k}]/N_k\} > (1/N_k^2)$, where $\hat{\sigma}^2_{\bar{\mu},N_k}$ is the mean square error of $\bar{\mu}_{N_k}$.*

Proof. The sample variance of the estimator $\hat{\xi}_k$ is $\mathrm{Samp\ Var}(\hat{\xi}) = [\sum_{i=0}^{N_k}(\hat{\xi}_i - \bar{\xi})^2]/(N_k - 1)$. From the Nth product to the $(N+1)$th product, the differential movement of this sample variance is

$$\begin{aligned}
&\frac{(\hat{\xi}_{N_k+1} - \bar{\xi})^2(N_k - 1) - \sum_{i=1}^{N_k}(\hat{\xi}_i - \bar{\xi})^2}{N_k(N_k - 1)} \\
&= \frac{(\hat{\mu}_{N_k+1} - \bar{\mu}_{ik})^2(N_k - 1) - \sum_{i=1}^{N_k}(\hat{\mu}_{ik} - \bar{\mu}_{N_k})^2}{N_k(N_k - 1)} \\
&= \frac{(\hat{\mu}_{N_k+1} - \bar{\mu}_{ik})^2 - \hat{\sigma}^2_{\bar{\mu},N_k}}{N_k}.
\end{aligned}$$

The agency's uncertainty increases only when this differential exceeds the expected reduction of posterior variance, which obeys $d/dN_k\,\mathrm{Var}^p[\hat{\xi}_k] > \Phi(\delta A - m, s)N_k^{-2}$. ■

Proof of Corollary 5.1

We consider the case where the product submitted by a bad firm ($N_{k^B} > 0, \hat{\xi}_{k^B} > 0$) and one submitted by an unknown firm ($N_{k^0} = 0 \Rightarrow \hat{\xi}_{k^0} = 0$) have identical danger and generate identical experimental histories for the danger variable X, or

$$\mu^{\xi>0}_{1,j,k^B} \equiv \mu^{\xi=0}_{2,j,k^0}$$

and

$$H\big(X_{1,j,k^B}(t) \mid \mu^{\xi>0}_{1,j,k^B}\big) \equiv H\big(X_{2,j,k^0}(t) \mid \mu^{\xi=0}_{2,j,k^0}\big), \quad \forall t \qquad \textbf{(A5)}$$

Note that this case implies that a "bad" firm has submitted a "better than expected" product. This is precisely the case we must consider, however, for the disadvantage of a bad firm reputation persists even when different firms submit substantively identical products.[24]

I first show that the disadvantage of a bad reputation ($\hat{\xi}_{k^B} > 0$) and the advantage of familiarity ($N_k > 0$) are both of order t^{-1}. By assumption $X(t)$ starts at zero, which by Eq. (2a) implies that $\hat{\mu}(t = 0) = m$ for an unknown firm ($N_{k^0} = 0$). Then for any $N_{k^B} > 0$ and positive $\hat{\xi}_{k^B}$, a Bayesian agency exploits the history of the firm's submissions by setting $\hat{\mu}^{\xi>0}_{i,k^B,t=0} = m + \hat{\xi}_{k^B}$, according to Eq. (7a). Now consider two products ($i = 1, 2$) submitted by two different firms (k^B and k^0). The products will still be treated differently in regulation, as

$$\hat{\mu}^{\xi>0}_t\big\{H\big(X_{1,j,k^B}(t) \mid \mu^{\xi>0}_{1,j,k^B}\big)\big\} > \hat{\mu}^{\xi=0}_t\big[H\big(X_{2,j,k^0}(t) \mid \mu^{\xi=0}_{2,j,k^0}\big)\big] \quad \forall t < \infty.$$

Even with identical data, then, the posterior danger estimates of the two products converge only in the asymptote because the priors differ. By identity of the histories (A5), $x_{1t} = x_{2t}, \forall t$, and the difference between the posterior estimates is a parametric (nonstochastic) function of t^{-1}:

$$\begin{aligned}
\hat{\mu}^{\xi>0}_t[H] - \hat{\mu}^{\xi=0}_t[H] &= \frac{(m + \xi)/s + x_{1t}/\sigma^2}{1/s + t/\sigma^2} - \frac{m/s + x_{2t}/\sigma^2}{1/s + t/\sigma^2} \\
&= \xi\sigma_i^2\big(\sigma_i^2 + st\big)^{-1} \qquad \textbf{(A6)}
\end{aligned}$$

Let the time that it would take for the product of an unknown firm to get approved be

$$t_{stop} = \inf\big[t,\ \text{s.t. } \hat{\mu}^{\xi=0}_{i,k^0,t} < \eta(t) \mid \mu^{\xi=0}_{i,k^0} < \delta A\big]$$

and similarly define

$$t_{app,k^B} = \inf\big[t,\ \text{s.t. } \hat{\mu}^{\xi>0}_{i,k^B,t} < \eta(t) \mid \mu^{\xi>0}_{i,k^B} < \delta A\big].$$

By Eqs. (A5) and (A6), $t^{\xi=0}_{stop} < t^{\xi>0}_{app,k^B}$. But then a sufficient condition for firm $k = k^B$ to receive quicker approval is

$$[\delta A - S_{k^B}(t_{stop})^2] - [\delta A - S_{k^0}(t_{stop})^2] \geq \hat{\mu}^{\xi>0}_{t_{stop},k^B}[H] - \hat{\mu}^{\xi=0}_{t_{stop},k^0}[H]$$

Notice that the comparison is conducted at $t \equiv t_{stop}$. From (A6), this occurs when

$$\hat{\xi}_{k^B} \leq \frac{\sigma_i^2 + st_{stop}}{\sigma_i^2}[S_{k^0}(t_{stop})^2 - S_{k^B}(t_{stop})^2]. \qquad \textbf{(A7)}$$

It can be shown that $S_k(t_{stop})^2$ has order t^{-2}_{stop}, hence the term in brackets will also have order t^{-2}_{stop}. Approximately, then, the right-hand side of (A7) has order t^{-1}_{stop}. Hence as t rises, (A7) is ever less likely to hold. Familiarity advantages hold for products that receive relatively early approval. □

The propositions for early-entrant protection invoke the agency taking into account its future decisions. Lemma 2 shows that the value function in Eq. (10) in the text satisfies a sequential optimality across submissions. It cannot be optimal, in other words, for the agency to game the sequential problem by holding off on the present drug to wait for other, perhaps superior drugs to come along.

[24] The demonstration here is direct and mostly algebraic. A more analytic proof is possible using results from stochastic calculus, which the author can supply upon request.

Lemma 2. *Let t^*_{app} be the optimal approval time under the payoff above. The policy of waiting for a short time $(t^*_{app}+\tau)$ so that another (possibly superior) drug may be approved, and the approval payoff reduced, can never be optimal.*

Proof. I consider the case where only one more product, the $(N+1)$th drug, is expected to be submitted in the future. The result here is generalizable to a finite or infinite sequence of product submissions. I temporarily relax the assumption that all determinants of the pipeline value are invariant and allow $G^*(t)$ to vary during the interval τ. For the $(N+1)$th drug, I consider three cases.

Case 1: If the $(N+1)$th drug is not submitted during the interval τ; then G^* is invariant over the interval. Let A_θ represent those determinants of the approval payoff that are not a part of the pipeline value. Then (suppressing i) by waiting for an interval τ, the agency gets a payoff of

$$A'_N = e^{-\delta\tau}\left\{A_\theta\left[1-e^{-\delta(E_\chi[t_{sub(N+1)}(\lambda)]+E_\chi[t^*_{app(N+1)}]-\tau)}c_j\,G^*_{N+1,\chi}(c_j)\right]\right\}.$$

By not waiting, the agency received

$$A_N = A_\theta\left[1-e^{-\delta(E_\chi[t_{sub(N+1)}(\lambda)]+E_\chi[t^*_{app(N+1)}])}c_j\,G^*_{N+1,\chi}(c_j)\right].$$

It is optimal to wait iff $A'_N > A_N$ or iff

$$\begin{aligned}
&e^{-\delta\tau}\left\{A_\theta\left[1-e^{-\delta(E_\chi[t_{sub(N+1)}(\lambda)]+E_\chi[t^*_{app(N+1)}]-\tau)}c_j\,G^*_{N+1,\chi}(c_j)\right]\right\}\\
&\quad> A_\theta\left[1-e^{-\delta[E_\chi[t_{sub(N+1)}(\lambda)]+E_\chi[t^*_{app(N+1)}]]}c_j\,G^*_{N+1,\chi}(c_j)\right]\\
&\quad\Rightarrow e^{-\delta\tau}-e^{-\delta(E_\chi[t_{sub(N+1)}(\lambda)]+E_\chi[t^*_{app(N+1)}])}\\
&\quad> 1-e^{-\delta(E_\chi[t_{sub(N+1)}(\lambda)]+E_\chi[t^*_{app(N+1)}])},
\end{aligned}$$

which is true only for $\tau<0$. □

Case 2: The $(N+1)$th drug is submitted during the interval τ. So G^* is not invariant. Let the approval distribution under this alternative policy be $G^{*\prime}$. Then waiting for τ is optimal iff

$$\begin{aligned}
&e^{-\delta\tau}\left[1-e^{-\delta(E_\chi[t_{sub(N+1)}(\lambda)]+E_\chi[t^*_{app(N+1)}-\tau])}\right]G^{*\prime}\\
&\quad>\left[1-e^{-\delta(E_\chi[t_{sub(N+1)}(\lambda)]+E_\chi[t^*_{app(N+1)}])}\right]G^*\\
&\quad\Rightarrow G^{*\prime}(\delta\tau)-G^*(\delta\tau)>1-e^{-\delta\tau} \qquad \textbf{(A8)}
\end{aligned}$$

But given the functional form of G^*, $G^{*\prime}(0)=0$, and $\lim_{t\to 0} g^{*\prime}(t)=0$, i.e., the cdf, pdf, and hazard are lowest at the beginning of review. So Eq. (A8) cannot be satisfied.

Case 3: The $(N+1)$th drug has already been submitted, but has not been approved, before t^*_{app}. Again, it is optimal to wait iff Eq. (A8) holds. Assume it is optimal to wait to approve the Nth drug for the interval τ. But then by the same policy it must be optimal to wait on the approval of the $(N+1)$th drug for an interval τ to see if the Nth drug gets approved. Then $\int_{t_{sub(N+1)}}^{t_{sub(N+1)}+\tau} g^{*\prime}\,dt=0$ for the interval, and Eq. (A8) cannot be satisfied. □

All propositions for early-entrant protection make use of the following lemma.

Lemma 3. *Given two drugs with identical danger ($\mu_N=\mu_{N+1}$), the absence of early-entrant protection for the Nth relative to the $(N+1)$th requires*

$$\gamma_{N,j}\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*c_{N+1,j}\right)\le\gamma_{N+1,j}(1-\gamma_{N,j}) \qquad \textbf{(A9)}$$

Proof. Early-entrant protection cannot hold if for any two drugs $i=N, N+1$, it is the case that $A_N\le A_{N+1}$, or

$$\begin{aligned}
&\psi_j L_j\gamma_{N,j}\prod_{i=1}^{N-1}(1-\gamma_{ij})\left[\prod_{i=N+1}^{N_{\max}}\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*c_{ij}\right)\right]\\
&\le\psi_j L_j\gamma_{N+1,j}\prod_{i=1}^{N}(1-\gamma_{ij})\left[\prod_{i=N+2}^{N_{\max}}\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*c_{ij}\right)\right]\\
&\Leftrightarrow\left\{\begin{array}{l}\psi_j L_j\gamma_{N,j}[(1-\gamma_{N-1,j})(1-\gamma_{N-2,j})L]\\ \quad\times\left[\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*c_{N+1,j}\right)\right.\\ \quad\times\left.\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*c_{N+2,j}\right)L\right]\end{array}\right\}\\
&\le\left\{\begin{array}{l}\psi_j L_j\gamma_{N+1,j}[(1-\gamma_{N,j})(1-\gamma_{N-1,j})L]\\ \quad\times\left[\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*c_{N+2,j}\right)\right.\\ \quad\times\left.\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*c_{N+3,j}\right)L\right]\end{array}\right\}
\end{aligned}$$

Cancel like terms in the telescoping products on both sides of the inequality to get (A9). □

Proof of Proposition 6

Strong early-entrant protection cannot hold if for any two drugs $i=N, N+1$, it is the case that $A_N\le A_{N+1}$, or (A9). Now assume stationarity of c_{ij}. Then in expectation, (A9) reduces to

$$c_{ij}\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*(\cdot)c_{ij}\right)\le c_{ij}(1-c_{ij}),$$

which cannot be satisfied because $\mathrm{P}[t^*_{app}=0]=0$. For concavity, notice that the likelihood of early-entrant protection is increasing in $e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*$. The first derivative of this expression with respect to N is negative, the second positive. □

Propositions 7 and onward require the following lemma.

Lemma 4. *Let $\gamma_N=(\alpha/\beta)$ (where $0<\alpha\le\beta$) and let the improvement for γ_{N+1} be represented as $\gamma_{N+1}=[(\alpha+\theta)/\beta]$, where $\theta\in(0,\beta-\alpha]$. Then conditioning on identical danger, the expected approval time for drug $i=N+1$ is less than or equal to that for drug $i=N$ whenever*

$$e^{-\delta(E_\chi[t_{sub}(\lambda)]+E_\chi[t^*_{app}])}G^*\ge\left(1-\frac{\theta\beta}{\alpha^2+\alpha}\right). \qquad \textbf{(A10)}$$

Proof. Rewrite (A9) as

$$\begin{aligned}
&\left(\frac{\alpha}{\beta}\right)\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*\left(\frac{\alpha+\theta}{\beta}\right)\right)\\
&\quad\le\left(\frac{\alpha+\theta}{\beta}\right)\left(1-\frac{\alpha}{\beta}\right)\\
&\quad\Rightarrow\frac{e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*[\alpha^2+\theta\alpha]}{\beta}\le\frac{\theta\beta-(\alpha^2+\theta\alpha)}{\beta}\\
&\quad\Rightarrow-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*\le\frac{\theta\beta}{\alpha^2+\theta\alpha}-1.
\end{aligned}$$

Proof of Proposition 7

Immediate. Unless $\theta>0$, (A10) cannot be satisfied.

Proposition 8 and its corollaries require the following result.

Lemma 5. *Given maximal concave stochastic curing improvement, and given two drugs with identical danger ($\mu_N = \mu_{N+1}$), then if there was early-entrant protection for the $(N-1)$th relative to the Nth, early-entrant protection must prevail for the Nth relative to the $(N+1)$th.*

Proof. Let ε be an infinitesimal positive quantity (the epsilon of analysis). For any three drugs $\gamma_N, \gamma_{N+1}, \gamma_{N+2}$, we may represent maximal concave stochastic improvement as

$$\gamma_N = \frac{\alpha}{\beta}, \quad \gamma_{N+1} = \frac{\alpha+\theta}{\beta}, \quad \gamma_{N+2} = \frac{(\alpha+2\theta-\varepsilon)}{\beta}. \qquad \textbf{(A11)}$$

Assuming that early-entrant protection prevailed between the *Nth* and the $(N+1)$*th* drug is equivalent to assuming

$$e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^* < 1 - \frac{\theta\beta}{\alpha^2+\theta\alpha}. \qquad \textbf{(A12)}$$

Then there is no early-entrant protection iff $A_{N+1} \le A_{N+2}$, or (suppressing j)

$$\psi L \gamma_{N+1} \prod_{i=1}^{N}(1-\gamma_i)\left[\prod_{i=N+2}^{N_{\max}}\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^* c_i\right)\right]$$
$$\le \psi L \gamma_{N+2} \prod_{i=1}^{N}(1-\gamma_i)\left[\prod_{i=N+3}^{N_{\max}}\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^* c_i\right)\right],$$

which by Lemma 3 implies

$$\gamma_{N+1}\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t_{app}*])}G^* c_{N+2}\right) \le \gamma_{N+2}(1-\gamma_{N+1}).$$

Using (A11), this can be expressed as

$$\left(\frac{\alpha+\theta}{\beta}\right)\left(1-e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*\left(\frac{\alpha+2\theta-\varepsilon}{\beta}\right)\right)$$
$$\le \left(\frac{\alpha+2\theta-\varepsilon}{\beta}\right)\left(1-\frac{\alpha+\theta}{\beta}\right)$$
$$\Rightarrow (\alpha+\theta)\left\{\beta - e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*[\alpha+2\theta-\varepsilon]\right\}$$
$$\le (\alpha+2\theta-\varepsilon)(\beta-(\alpha+\theta))$$
$$\Rightarrow e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^* \ge 1 - \frac{\beta(\theta-\varepsilon)}{(\alpha+\theta)(\alpha+2\theta)-\varepsilon(\alpha+\theta)}.$$

Now by (A12) and Lemma 4,

$$1-\frac{\theta\beta}{\alpha^2+\theta\alpha} \ge e^{-\delta(E_\chi[t_{sub}]+E_\chi[t^*_{app}])}G^*$$
$$\ge 1 - \frac{\beta(\theta-\varepsilon)}{(\alpha+\theta)(\alpha+2\theta)-\varepsilon(\alpha+\theta)}.$$

So the absence of early-entrant protection now requires

$$1-\frac{\theta\beta}{\alpha(\alpha+\theta)} \ge 1 - \frac{\beta(\theta-\varepsilon)}{(\alpha+\theta)(\alpha+2\theta)-\varepsilon(\alpha+\theta)}$$
$$\Rightarrow \frac{\theta}{\alpha} \le \frac{\theta-\varepsilon}{\alpha+2\theta-\varepsilon}.$$

But $\lim_{\varepsilon\to 0}[(\theta-\varepsilon)/(\alpha+2\theta-\varepsilon)] = [\theta/(\alpha+2\theta)$: contradiction. □

Proof of Proposition 8

Let any γ_1 and γ_2 satisfy (A11), in which case (1) $dE[\gamma]/dN_{i,j=J} > 0$ and (2) $d^2E[\gamma]/dN^2_{i,j=J} < 0$ are satisfied for all $N_{i,j=J}$. Then (1) follows immediately from Lemma 5. For (2), note that whether or not there is improvement, $\lim_{N\to\infty} G^*(N,t) = 0$. Then for any sequence $i = 1, \ldots, N, \ldots, N_{max}$, it must be the case that for some two drugs $i = N$, $N+1$, $A_N \le A_{N+1}$. Then by Lemma 5, weak early-entrant protection follows ever thereafter. □

Proof of Corollary 8.1

Note that in (A10) the right-hand side is strictly decreasing in α. But α is simply the numerator of γ_1. □

Proof of Corollary 8.2

The left-hand side of (A10) is strictly decreasing in λ_j. □

Proof of Corollary 8.3

Assume that each firm k develops N_k drugs randomly for a set of J_N diseases and that, for any disease $j = J$, N_J drugs are submitted and only one receives strong early-entrant protection. Assume also that the order of submission for any K firms is random (this is unlikely in a strategic environment). Then for any firm, the likelihood of receiving strong early-entrant protection for at least one drug is equal to $[1/N_j] \times [N_k/J_N]$. □

REFERENCES

Arnold, R. Douglas. 1990. *The Logic of Congressional Action*. New Haven, CT: Yale University Press.

Bartel, Ann, and Lacy Glenn Thomas. 1987. "Predation through Regulation." *Journal of Law and Economics* 30 (2): 239–64.

Becker, Gary S. 1985. "Public Policies, Pressure Groups, and Dead Weight Costs." *Journal of Public Economics* 28: 330–47.

Bendor, Jonathan, Dilip Mookherjee, and Debraj Ray. 2001. "Aspiration-Based Reinforcement Learning in Repeated Games: An Overview." *International Game Theory Review* 3 (2–3): 159–74.

Bernstein, Marver H. 1955. *Regulating Business by Independent Commission*. Princeton, NJ: Princeton University Press.

Brehm, John, and Scott Gates. 1997. *Working, Shirking and Sabotage*. Ann Arbor: University of Michigan Press.

Breyer, Stephen. 1982. *Regulation and Its Reform*. Cambridge, MA: Harvard University Press.

Carpenter, Daniel P. 2001. *The Forging of Bureaucratic Autonomy: Reputations, Networks and Policy Innovation in Executive Agencies, 1862–1928*. Princeton, NJ: Princeton University Press.

Carpenter, Daniel P. 2002. "Groups, the Media, Agency Waiting Costs, and FDA Drug Approval." *American Journal of Political Science* 46 (3): 490–505.

Carpenter, Daniel P., and Michael M. Ting. 2004. "A Theory of Approval Regulation with Endogenous Submissions." Manuscript. Harvard University.

Chernoff, Herman. 1968. "Optimal Stochastic Control." *Sankhya — The Indian Journal of Statistics* Series A 30 (3): 221–252.

Chhikara, R. S., and J. Leroy Folks. 1989. *The Inverse Gaussian Distribution: Theory, Methodology, and Applications*. New York: Marcel Dekker.

Crandall, Robert W., and Kenneth Flamm, eds. 1989. *Changing the Rules: Technological Change, International Competition, and Regulation in Communications*. Washington, DC: Brookings Institution.

DeGroot, Morris H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill.

DiMasi, Joseph A. 1996. "A New Look at United States Drug Development and Approval Times." *American Journal of Therapeutics* 3: 1–11.
DiMasi, Joseph A., Ronald W. Hansen, Henry G. Grabowski, and Louis Lasagna. 1991. "Cost of Innovation in the Pharmaceutical Industry." *Journal of Health Economics* 10 (1991): 107–42.
Dixit, Avinash K. 1993. *The Art of Smooth Pasting*. Chur. Switzerland: Harwood Academic.
Dranove, David, and David Meltzer. 1994. "Do Important Drugs Reach the Market Sooner?" *RAND Journal of Economics* 25 (3): 402–23.
Epstein, Steven. 1996. *Impure Science: AIDS, Activism, and the Politics of Knowledge*. Berkeley: University of California Press.
General Accounting Office ("GAO"). 1995. *FDA Review Time for Drugs has Decreased in Recent Years*. Washington, DC: GAO.
Gordon, Sanford. 1999. "*Managing Fairness*." Ph.D. Dissertation. Princeton University.
Grabowski, Henry G., and John M. Vernon. 1976. "Structural Effects of Regulation on Innovation in the Ethical Drug Industry." In *Essays on Industrial Organization in Honor of Joe S. Bain*; eds. Robert T. Masson and P. David Qualls. Cambridge, MA: Lippincott, Ballinger, 1976, 181–205.
Grabowski, Henry G., and John M. Vernon. 1983. *The Regulation of Pharmaceuticals*. Washington, DC: AEI Press.
Hammond, Thomas H., and Paul Thomas. 1989. "The Impossibility of a Neutral Hierarchy." *Journal of Law, Economics and Organization*, 5 (1): 155–84.
Heimann, C. F. Larry. 1997. *Acceptable Risks: Politics, Policy and Risky Technologies*. Ann Arbor: University of Michigan Press.
Hilts, Philip. 2003. *Protecting America's Health: The FDA, Business, and One Hundred Years of Regulation*. New York: Knopf.
Hinich, Melvin J., and Richard Staelin. 1980. *Consumer Protection Regulation and the U.S. Food Industry*. New York: Pergamon Press.
Huntington, Samuel P. 1953. "The Marasmus of the ICC: The Commission, the Railroads and the Public Interest." *Yale Law Journal* 61: 467–509.
Jackson, Charles O. 1970. *Food and Drug Legislation in the New Deal*. Princeton, NJ: Princeton University Press.
Joskow, Paul. 1974. "Inflation and Environmental Concern: Structural Change in the Process of Public Utility Price Regulation." *Journal of Law and Economics* 17: 291–327.
Jovanovic, Boyan. 1979. "Job Matching and the Theory of Turnover." *Journal of Political Economy* 87 (5): 972–90.
Karatzas, Ioannis, and Steven E. Shreve. 1991. *Brownian Motion and Stochastic Calculus*. 2nd ed. New York: Springer-Verlag.
Kyle, Margaret. 2002. "Entry in Pharmaceutical Markets." Mimeo. Carnegie–Mellon.
Miroschnichenko, T. P. 1975. "Optimal Stopping of an Integral of a Wiener Process." *Theory of Probability and Its Applications* 9 (4): 355–92.
Mullainathan, Sendhil. 2002. "A Memory-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 117 (3): 735–74.
Noll, Roger G. 1973. *Economic Aspects of Television Regulation*. Washington, DC: Brookings Institution.
Noll, Roger G. 1985. "Government Regulatory Behavior: A Multidisciplinary Survey and Synthesis." In *Regulatory Policy and the Social Sciences*, ed. R. Noll. Berkeley: University of California Press.
Olson, Mary. 1995. "Substitution in Regulatory Agencies: FDA Enforcement Alternatives." *Journal of Law Economics and Organization* 12 (2): 376–407.
Olson, Mary. 1997. "Firm Characteristics and the Speed of FDA Approval." *Journal of Economics and Management Strategy* 6 (2): 377–401.
Peltzman, Sam. 1973. "An Evaluation of Consumer Protection Legislation: The 1962 Drug Amendments." *Journal of Political Economy* 81 (5): 1049–91.
Peltzman, Sam. 1976. "Toward a More General Theory of Regulation." *Journal of Law and Economics* 19: 211–40.
Quirk, Paul. 1980. "The Food and Drug Administration." In *The Politics of Regulation*, ed. James Q. Wilson. New York: Basic Books.
Rabin, Matthew, and Joel L. Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *Quarterly Journal of Economics* 114 (February): 37–82.
Rothenberg, Lawrence S. 1994. *Regulation, Organizations, and Politics: Motor Freight Policy at the Interstate Commerce Commission*. Ann Arbor: University of Michigan Press.
Schultz, Kenneth. 1998. "Domestic Opposition and Signaling in Crisis Bargaining Games." *American Political Science Review* 92 (4): 829–44.
Shepp, Larry. 1969. "Explicit Solutions to Some Problems of Optimal Stopping." *Annals of Mathematical Statistics* 40: 993–1010.
Spence, David B. 1999. "Managing Delegation Ex Ante: Using Law to Steer Administrative Agencies." *Journal of Legal Studies* 28 (2): 413–60.
Stigler, George J. 1975. *The Citizen and the State: Essays on Regulation*. Chicago: University of Chicago Press.
Taylor, Howard M., and Samuel Karlin. 1998. *An Introduction to Stochastic Modeling*. 3rd ed. New York: Academic Press.
Thomas, Lacy Glenn. 1990. "Regulation and Firm Size: FDA Impacts on Innovation." *RAND Journal of Economics* 21 (4): 497–517.
Ting, Michael M. 2003. "Flying Blind: A Theory of Bureaucratic Learning." Manuscript, Columbia University.
Viscusi, W. Kip. 1992. *Fatal Tradeoffs*. New York: Oxford University Press, 177.
Vogel, David. 1996. *Kindred Strangers: The Uneasy Relationship between Politics and Business in America*. Princeton, NJ: Princeton University Press, chap. 8.
Wilson, James Q. 1980. *The Politics of Regulation*. New York: Basic Books.