# The Downside of Deadlines

Daniel Carpenter and Justin Grimmer

February 7, 2009

**Abstract**

Deadlines are ubiquitous institutions in government decision making, constraining both agencies and courts. Yet these institutions are almost entirely ignored in formal models in institutional political science. We analyze deadlines as exogenously imposed institutions upon a government decision maker, as a means of elected officials exercising control over the the duration of administrative decision processes. Our formal model demonstrates how deadlines are successful at lowering the time to administrative decision. Yet our analysis also illuminates the unappreciated pitfalls of deadlines, or their "downside." The effect of deadlines on regulatory behavior is highly non-linear, making imposition of deadlines a difficult task for even highly rational agents. Further, our model predicts that deadlines will increase the variance in the review time distribution under a large set of conditions and predicts that deadlines will increase the error rate in regulatory behavior, often in an exponential fashion. Our formal analysis helps to explain an expanding set of empirical findings about the effects of deadlines and suggests some of the limits of deadlines as an effective tool of control over policymaking and bureaucratic decisions.

# 1 Introduction

Because politicians or bureaucratic superiors wish to limit administrative and regulatory delay, they often impose deadlines of one form or another upon agencies and other government actors. Deadlines for decision making are a common feature of the judicial realm (Abbott, 1987), and in recent decades deadlines have also governed administrative rulemaking (Gersen and O'Connell, 2008) as well as regulatory decisionmaking in numerous areas (Gersen and O'Connell, 2008; O'Connell, 2008). Recent empirical work, moreover, has demonstrated that deadlines influence the timing and duration of regulatory decisions (Carpenter, Zucker and Avorn, 2008; Gersen and O'Connell, 2008; Yackee and Yackee, 2008), and has also presented evidence of an association between just-before-deadline approval decisions and subsequent drug safety problems in the arena of pharmaceutical regulation (Carpenter, Zucker and Avorn, 2008).

In this article we provide a theoretical explanation for these relationships, and we point theoretically to several features of deadlines that analysts have not considered. To analyze the effects of deadline bonuses on government decision making, we develop a model of case review by a dynamically rational but uncertain agent. A deadline policy is conceived as a time for decision combined with a bonus for meeting it (or a penalty for missing it). For example, this is the policy used to govern drug review by regulatory agencies in the United States and European Union. Some of the results of the model are expected and not entirely surprising. We show that deadline bonuses are likely to decrease the review time of cases, and that the regulator's error rate (the tendency to approve cases that should not have been approved) is increasing in the size of bonus offered.

Yet analysis of our model generates several results that run counter to popular intuition and previous research. We show, for instance, that deadlines can increase the variance of government decision making even when their mean effect is one of acceleration. This is not an isolated or knife-edge occurrence but occurs under a wide range of plausible parameter values for the model. In other words, deadlines can make government decisions *less* predictable, not more so. Second, the induced error rate from deadlines may be nonlinearly (exponentially) increasing in the size of the bonus offered for meeting the deadline, in ways that are highly complex and probably both unobservable and unpredictable. These considerations suggest that deadlines are quite possibly an effective but problematic tool for political control; political authorities may not know exactly "what they are getting" or "what they are getting into" when they impose deadlines upon administrative agents. This notion has implications for further theoretical and empirical study of govern-

ment decision making and bureaucratic politics, and is an issue to which we return in the conclusion.

Our analysis also has broad methodological implications for the empirical analysis of government decision making and regulatory review. Specifically, we show that deadlines cause a discontinuity in the hazard rate of product approval. This discontinuity causes large biases in the estimates of coefficients if parametric or semi-parametric methods are applied to the review process (as in Carpenter 2002; Whitford 2005; Kosnik 2006). Hence some of the more commonly studied duration processes in political science and related social sciences may require re-examination through more nuanced statistical models that include the presence of deadlines, even if the deadlines are "soft" or only partially binding as normative guidelines to decision making.

As with all such theoretical efforts, our model has some basic limitations, most of which follow from the fact that the model is developed for a continuous-time, continuous state, decision-theoretic context. Hence the first limitation of our model is that it is not game-theoretic and does not analyze the possibility of continuous interaction between the government decision maker (as an agent) and the elective institution (say Congress or the President, as a principal) over time. Such an analysis might reveal important learning dynamics and adaptation by both players. Yet one vital response to this objection is that if decision-theoretic analysis of modern government decision making yields rather daunting complexities – and we demonstrate below that it does – then fully strategic equilibrium behavior in a game-theoretic world is likely (though perhaps not certain) to add even more. Second, the sort of complex, non-linear effects we show to be associated with deadlines are unlikely to be fully anticipated by either principals or agents themselves in a strategic setting. Indeed, the modeling of agents as fully dynamically rational and both attentive and faithful to continuation values is asking much of any human agent (Berns, Laibson and Loewenstein, 2007); our model adds a realistic but complex stochastic process along with institutional context. Finally, if a fully strategic context were to be adopted, then the analysis would be vastly more complex and quite likely unresolvable analytically; at the very least, it would need to be conducted in the still-evolving and limited field of differential games.

## 2 Deadlines in Administrative Practice and in Scholarly Analysis

The prospect that government decision makers can delay implementation or choice makes the problem of political control all the more challenging. Not surprisingly, the duration of government decision making has become a source of increasing scholarly analysis (Carpenter, 2002; Whitford, 2005; Kosnik, 2006). A

common institutional procedure adopted to address the problem of administrative delay is the imposition of a deadline upon government agents. Deadlines – rules that specify a maximal length of time for a decision or case review to be taken and an associated reward for meeting the maximal time (or penalty for missing it) – would appear to be widespread in institutional politics.[1] Congress uses deadlines and statutory hammers to produce quicker action in a wide variety of delegation decisions (Abbott, 1987), including environmental regulation (Morgenstern, 1993) and national traffic safety regulation (Mashaw and Harfst, 1990, 69-83). The Environmental Protection Agency (EPA), for instance, is given a set amount of time to issue standards for certain toxic substances or pollutants under its enabling legislation, and it is also charged with revisiting its decisions every five to seven years. The 1998 amendments to the Federal Insecticide, Fungicide and Rodenticide Act to re-register 700 substances within a period of ten years (each registration and re-registration is separate, costly and time-consuming) (Morgenstern, 1993, 245). In recent years, congressional statutes have charged the FDA with meeting deadlines of six months for the review of more therapeutically vital "priority" drugs and (since 1997) ten months for most other drug reviews. Other health regulators worldwide – including the European Medicines Evaluation Agency (EMEA) – have adopted similar regulatory review deadlines.

The emerging literature on deadlines in government decision making is small but has generally found that deadlines do indeed accelerate administrative decision processes. Two recent comprehensive analyses of deadlines in the study of rulemaking by U.S. government agencies, by Gersen and O'Connell (2008) and by Yackee and Yackee (2008), have both produced general evidence for the proposition that deadlines induce quicker decisions by federal administrators charged with rulemaking. And a recent study of deadlines in drug regulation (Carpenter, Zucker and Avorn, 2008) shows evidence for acceleration of FDA review times and for increased error rates associated with the drug review deadlines.

In the social sciences more generally, previous formal and experimental studies of deadlines have been conducted largely in the setting of negotiation (e.g., Fershtman and Seidmann 1993; Stuhlmacher and Champagne 2000; Gneezy et al. 2003; Moore 2004). The analysis of delay in formal models of negotiation is useful but also compromised by the fact that most such models focus on stationary equilibria where the his-

[1] In some cases, deadline institutions impose a more smoothed penalty or reward structure, such as the practice of a one-third or one-half-grade reduction "for every day late" used in the grading of papers or assignments at schools and universities. Because our analysis focuses upon political and administrative decisions, where discrete payment structures are used most commonly, we do not explore smoothed bonus payments in this model. Discrete payment/penalty structures correspond to the form of deadlines used in American and European regulation of drugs, medical devices and food additives, as well as to deadline structures in administrative rulemaking and judicial arenas.

tory of previous bargaining does not matter to present bargaining.[2] Where formal models have relaxed the stationarity assumption, they have assumed complete information or have incorporated a very limited notion of uncertainty (e.g., Fershtman and Seidmann 1993, 307, 309). The same can be said for many experimental studies of deadlines, which either focus on negotiation or incorporate very limited structures of uncertainty and information (e.g., Moore 2004). Since much of government decision making involves *learning* about an uncertain environment, these models sacrifice what is perhaps the major focus of interest for students of institutions and policy.

It is the combination of learning over time, uncertainty and deadlines, then, that is of interest to so many students of political decision making, administrative law and politics, and public policy. Yet the combination of (a) general uncertainty with both (b) dynamic choice and (c) deadlines seems to have eluded both the world of theory and the world of empirics. We now turn to elaborate a model which contains all three elements.

# 3 A Generalized Dynamic Model of Government Approval

The model analyzed in this paper builds upon and generalizes prior literature on continuous-time learning in government decision-making. The agent is faced with a learning problem, where the quality and danger of a product or "case" submitted for judgment or approval are unknown quantities (Carpenter, 2004, 2002). The regulator wants to approve a case when the stochastic payoff from doing so outweighs the (1) the stochastic danger of the case and (2) the value of waiting for more information.

## 3.1 Incentives for Irreversibility, Including Reputation Protection

Our model presumes that the government decision maker perceives approval as irreversible or reversible only at significant cost. In the drug approval case, for instance, the regulator's primary objective when deciding to approve a new drug may be to protect her reputation (Carpenter, 2004). A reputation-protecting pharmaceutical regulator may wish to ensure that the rate of adverse safety events associated with a new drug is not materially different from the rate of adverse events associated with other drugs in the same therapeutic class. An agency promulgating a new rule may worry that its decision will be associated with highly publicized mistakes or disasters later on, problems that could have been anticipated by "signals" in

[2] In other words, the customary restriction to examining stationary equilibria means that the set of offers that players can either make or accept in any subgame is independent of the history of previous (rejected) offers which generate the status quo subgame; see Fershtman and Seidmann 1993.

the agency's evidentiary record. In these cases, the agent's task is to disentangle the information about a case's quality from the information provided about its danger, and to process both "quality" signals and "danger" signals while not suboptimally delaying the decision.

## 3.2 Notation and Assumptions

While generally conceived, one can use the running example of approval regulation of pharmaceuticals to make the following model more concrete. Under this specific rendering, "cases" may be construed as "drugs," "quality" as "efficacy", and "safety" as the absence of health-related hazards associated with the case. Let the cases (drugs) under review be indexed by $i$, let the review time for case $i$ be given by $t^i$ with the time of approval given by $t^i_{\text{app}}$. We suppose that each case is characterized by two parameters: its *quality* (in the drug example, its effectiveness in treating the disease) and its *danger*, or the rate at which that adverse events occur in association with the case. Suppose that a case's quality is a draw from a normal distribution, $\mu_i \sim \mathcal{N}(m, s)$ where $\mathcal{N}(m, s)$ represents the normal distribution with mean $m$ and variance $s$. The actual value of $\mu_i$ is unknown to the regulator, but is learned during the review process. The danger parameter is $\lambda_i \sim \text{Gamma}(\alpha, \beta)$.[3] The evidence process the regulator observes during the review of cases is an additive combination of continuous evidence, which we represent as Brownian motion with drift, and a compensated Poisson process–representing the discrete occurrence of adverse events.

### 3.2.1 Continuous Time Evidence of quality

The regulator collects continuous time evidence about a case's quality according to Brownian motion with drift, where the drift is determined by the (unobserved) quality of the case. Formally, the regulator observes

$$\text{cont}_i(t, \mu_i, \sigma_i) \quad = \quad \mu_i t + \sigma_i w(t) \tag{3.1}$$

where $w(t)$ is a standard normal distribution with mean zero and variance $t$. As we show below, the parameter $\sigma$ encodes the amount of information Equation 3.1 contains for the regulator: if $\sigma_i = 0$ then the regulator can immediately infer the quality of the case by examining the slope of Equation 3.1 and as $\sigma_i \to \infty$ Equation 3.1 contains no information about a case's quality.

[3]The danger parameter is not assumed known to the regulator at the start of the review process, but it is assumed that the regulator can form a "baseline" estimate of expected danger (the compensator term in the model below) and can compute a subjective (not necessarily Bayesian) expectation over the value of danger signals received.

### 3.2.2 Rare Occurrence of Adverse Events

The regulator also weighs the accumulation of danger signals against a baseline rate that serves as an expected rate of danger. We represent this balancing with a compensated, compound Poisson process, where the jumps occur in only one direction. In this model the jumps are all downward because the adverse events decrease the value of the case. Suppose that arrival time of events is exponentially distributed, with parameter $\lambda_i$. This implies that the number of adverse events that have occurred by time $t$, $J(t)$, is Poisson distributed with rate parameter $\lambda_i t$. Conditional on an event occurring, we suppose that the size of the jump is a draw from some distribution $G(Z)$ and suppose that the expected size of this jump is given by $\xi = \int_{\Re_+} Z dG(Z)$. If $G(Z)$ is degenerate, placing all probability on one value of $Z$, then the adverse events arrive according to a standard Poisson process.[4] Formally, the compensated compound Poisson process is,

$$\text{pois}_i(t, \lambda_i, J(t), \xi) \quad = \quad \lambda_i \xi t - \sum_{k=1}^{J(t)} Z_k \tag{3.2}$$

### 3.2.3 Defining the Evidence Process

The regulator observes an additive combination of Equations 3.1 and 3.2,

$$\begin{aligned} X(t) &= \text{cont}_i(t, \mu_i, \sigma_i) + \text{pois}_i(t, \lambda_i, J(t), \xi) \\ &= \mu_i t + \sigma_i w(t) + \lambda_i \xi t - \sum_{k=1}^{J(t)} Z_k \end{aligned} \tag{3.3}$$

Equation 3.3 is an example of a *Levy Process*: a stochastic process that combines Brownian motion with Poisson processes. This framework provides a general model for analyzing regulatory behavior, including as a special case the models in Carpenter (2002) and Carpenter (2004).[5] We will collect the path of $X(t)$ during the regulatory process into the history $H_i(t)$.[6]

## 3.3 Estimating quality from Evidence

Given that the regulator only observes $X(t)$ we first prove that the learning problem is identified: the regulator is able to disentangle the contribution of the quality of the case to $X(t)$ from the value contributed

[4] Suppose that $G(Z)$ has support for $Z > 0$. Further, suppose that the first moment exists. Our proofs require no other assumptions

[5] The process is pure Brownian motion with drift if $\lambda_i = 0$. Similarly, the process is only a compensated Poisson process if $\mu_i = 0$ and $\sigma_i = 0$

[6] In the appendix we offer a formal statement of the model developed here.

by Equation 3.2. Lemma 1 shows that this identification is possible and provides sufficient statistics for the learning problem.

**Lemma 1** (Identification of Learning Problem and Sufficient Statistics)**.** *Without loss of generality, for any $X_i(t)$, the history of $X_i(t)$, $\mathbf{H}_i(t)$ can be broken into three components: (1) the history of the continuous component due to Brownian motion with drift, $\mathbf{H}_i^B$, (2) the contribution of the compensation of the Poisson process $\mathbf{H}_i^c$, and (3) the history of the jumps $\mathbf{H}_i^J$. A sufficient statistic for $\mathbf{H}_i^B$ is the dual $(t, X_i(t)^*)$, where $X_i(t)^* = X_i(t) - \lambda\xi t + \sum_{k=1}^{J(t)} Z_k$, a sufficient statistic for $\mathbf{H}_i^J$ is $(t, \sum_{k=1}^{J_t} Z_k)$ and a sufficient statistic for $\mathbf{H}_i^c$ is (t, $\lambda_i\xi$)*

**Proof** *Proofs of all lemmata and propositions appear in the Appendix.*

The regulator can separate the components of $X_i(t)$ because the compensator for the adverse events $\lambda_i\xi t$ is known and the adverse events occur instantly (on a set of measure zero). During the time when these events cause the evidence process to jump downwards, the Brownian motion does not move, allowing the agent to identify movement due to the jumps and movement due to Brownian motion.

Lemma 1 is more than just a technical feature of our model, it captures an important component of regulatory dynamics. The rearrangement allows the agent to separate discrete and continuous evidence–much as a physician could separate more continuous outcome measures such as monthly pain or hypertension measurements from discrete events such as a myocardial infarction or an event that flagged severe hepatotoxicity (Olson, 1997; Carpenter, 2002). In the review of a dam licensing project by an agency like the Federal Energy Regulatory Commission (FERC) in the United States (Kosnik, 2006; Spence, 1999), a regulator might separate more continuous measures such as megawatt generation from more discrete outcomes such as failures or environmental catastrophes. Alternatively, the review of data from nuclear power plants, as part of an inspection or licensing operation, could include more continuous measurements for energy generation and rarer, discrete measurements for safety issues (Gordon and Hafer, 2005).

Using Lemma 1 we are able to define the regulator's posterior estimate of case quality, $\hat{\mu}_i(t)$, its variance $V(t)$, and the contribution of the compensated Poisson process to $X(t)$, $P(t)$. To define these quantities, we first calculate $X^*(t)$ by subtracting the components of $X(t)$ not related to case quality,

$X^*(t) = X(t) - \lambda\xi t + \sum_{k=1}^{J(t)} Z_k$ and (with a slight abuse of notation) call $x = X^*(t)$. Then,

$$\text{Posterior Mean} \equiv \text{E}(\mu_i|x) = \hat{\mu}_i(t) = \frac{(m/s) + (x/\sigma_i^2)}{(1/s) + (t/\sigma_i^2)} \tag{3.4}$$

$$\text{Var}\hat{\mu}_i \equiv \text{V}(t) = \frac{1}{(1/s) + (t/\sigma_i^2)} \tag{3.5}$$

$$\text{Poisson Process, Time t} \equiv \text{P}(t) = \lambda_i\xi t - \sum_{k=1}^{J(t)} Z_k \tag{3.6}$$

The posterior variance $V(t)$ incorporates the regulator's uncertainty about the quality of the case, while $P(t)$ reflects the value from the compensated Poisson process at time $t$. If $P(t) > 0$ then there have been fewer adverse events than expected, while if $P(t) < 0$, then there have been more adverse events than expected.

### 3.4 Filtered Evidence and Value Functions

The agent seeks to define an optimal stopping rule for the *filtered evidence process* found by combining Equations 3.4 and 3.6, $\hat{\gamma}(t)_i$[7],

$$\hat{\gamma}(t)_i = \hat{\mu}_i + \lambda_i\xi t - \sum_{k=1}^{J(t)} Z_k. \tag{3.7}$$

Suppose that there is some convex function $\hat{\gamma}(t)_i \times t \mapsto F(\hat{\gamma}(t)_i, t)$, that is twice differentiable with respect to both $\hat{\gamma}(t)_i$ and $t$. This function is a map from the current state of the filtered evidence process and time to the value experienced by the regulator.

### 3.5 Regulatory Objective

The regulator's objective is to define an optimal rule to stop $\hat{\gamma}(t)_i$ in order to maximize her payoff. Formally this implies,

$$\begin{aligned} &\max \text{E}e^{-\delta(t_{app})} \left\{ A + \text{E}_{\hat{\mu},\sum Z,t} \int_t^{\infty} e^{-\delta(y-t)} \left[ \mu(s,\omega) - \sum_{k=1}^{J_t} Z_k + \lambda_{i,j}\xi t \right] dy \right\} \\ &= \text{E}e^{-\delta(t_{app})} \left\{ A + \delta^{-1} \left( \mu^* \left[ t_{app}, \omega \right] - \sum_{k=1}^{J_t^*} Z_k + \lambda_{i,j}\xi t^* \right) \right\} \end{aligned} \tag{3.8}$$

[7]We say that this is the filtered evidence process because the Bayesian estimates of quality identifies the components of $X(t)$ necessary for estimating the quality–or "filters" the histories to measure the crucial components.

where $\delta$ is a discount factor, $A$ is an approval payoff which is static, positive and known with certainty throughout the review.[8] The quantities $\mu^*$, $t^*$, and $J^*(t)$ are the case's quality, time of approval and the number of jumps, respectively, at the optimal stopping time, and $y$ is a variable of integration.

### 3.6 The Optimal Stopping Rule and Its Properties

The regulator's optimal policy will be to observe the first passage of the $\hat{\gamma}(t)_i$ through a border. Upon approval of any product or case, the values $\mu_i$ are fully revealed as well as the difference between the jump process and the compensator–$\lambda_i \xi t^* - \sum_{k=1}^{J_t^*} Z_k$ and "payoffs" are realized. Proposition 1 states the optimal rule for stopping $\hat{\gamma}(t)_i$.

**Proposition 1.** *The case is approved when and only when, and if and only if, $\hat{\gamma}(t)_i$ passes for the first time through the following barrier,*

$$\begin{aligned} \eta^*(t) &= \delta(\lambda\xi - A) + \frac{1}{2\sigma^2} V(t) F_{\hat{\mu},\hat{\mu}}(\hat{\gamma}(t)_i, t) \\ &\quad + \lambda_i \int_{R_+} \left[ F\left(\hat{\gamma}(t)_i - Z, t\right) - F\left(\hat{\gamma}(t)_i, t\right) \right] dG(Z) \end{aligned} \tag{3.9}$$

*where $F_{\hat{\mu}\hat{\mu}}(\hat{\gamma}(t)_i, t)$ is the second partial derivative of the value function $F$ with respect to the filtered state variable $\hat{\mu}$, given a realization of $\hat{\mu}$ at time $t$.*

This border represents the optimal tradeoff between delaying the approval of a case and approving instantly. If the regulator delays the approval of a case, she receives more information, reducing the value of $V(t)$ and receives further compensation for the possible occurrence of adverse events $\delta\lambda\xi$. Approval provides the regulator with the payoff $\delta A$, while also allowing the regulator to avoid the risk of another adverse event in the next instance, which is represented by the (negative) value $\lambda_i \int_{R_+} [F(\hat{\gamma}(t)_i - Z, t) - F(\hat{\gamma}(t)_i, t)] \, dG(Z)$.

Figure 1, below, visualizes the filtered evidence process and the optimal stopping barrier from Proposition 1. The horizontal axis represents time, the vertical axis is the utility provided to the regulator, the red-line is the filtered evidence process for one product review, and the purple line represents the optimal stopping barrier. Adverse events appear as a discontinuous jumps downwards, reflecting the regulators' aversion to approving cases when there is an indication that a drug may have an exceptionally poor safety record for public use. The border slopes downward, as the value of more information decreases over the course of the regulatory history. A case is approved only if its evidence crosses the boundary, which occurs at the right-hand side of Figure 1.

[8] Carpenter (2004) provides intuition about the factors that determine this payoff.

Figure 1 About Here.

# 4 Deadlines and Regulatory Behavior

Many government agencies, including regulatory bodies, face time constraints of different sorts, some imposed from without and some induced by the structural characteristics of their tasks. In particular, politicians and courts concerned about limiting the delay associated with regulatory processes (whether governmental or private) may impose deadlines for decision making upon the agency (Gersen and O'Connell, 2008). In the case of pharmaceutical regulation in the United States, Congress has imposed "review-time goals" upon the U.S. Food and Drug Administration's Center for Drug Evaluation and Research (CDER), such that it is now expected to act upon ninety percent of all "standard" new product applications within 10 months or less (Carpenter, Zucker and Avorn, 2008). In other settings, deadlines are used to constrain the behavior of licensing agencies, product review agencies outside of the United States, and other forms of decision for regulatory agencies.

There are many possible rationales for deadlines, and all of them are exogenous to the model elaborated here. One benefit of deadlines might be that the regulator values time – for example, discounts the future – in a way markedly different from the way that citizens and their representatives do.

We now modify the formal model developed in Section 3 to generate predictions about the effects of exogenously imposed institutions on regulatory behavior. The deadline is an "institution" in the sense that it is an "if-then" rule: if the agent decades before the deadline, she receives the bonus, and forfeits the bonus if she surpasses it. The model shows that providing regulators with a bonus for approving cases before a deadline does result in more rapid approval. But, deadline bonuses can also make regulatory behavior *less predictable* and the effect of altering the deadline bonus on regulatory behavior is highly non-linear.

# 5 Deadlines Reduce Time To Approval

We consider flexible deadlines in the form of a bonus payment if the case is approved by the deadline. Imagine a deadline bonus $D$ ($D \in [0, \infty)$), which is awarded with certainty to the regulator if and only if she approves the case by an exogenous deadline $t^D$, s.t. $0 < t^D < \infty$. The finiteness of the deadline bonus means that the agency could allow some cases to endure past the deadline, depending upon specific values or evidence encountered in the case. Because the deadline bonus is imposed exogenously, it has a straightforward effect on the regulator's solution to the regulator's optimal stopping problem. The barrier specified in Equation 3.9 now takes one of two forms, depending on whether the deadline has elapsed.

Before the deadline has been reached, the regulator's adjusted stopping barrier is

$$\begin{aligned}\eta^{*} &= \frac{1}{2\sigma^{2}}V(t)F_{\hat{\mu},\hat{\mu}} - \delta(A + D - \lambda_{i}\xi) \\ &\quad +\lambda_{i}\int_{R_{+}}\left[F(\hat{\gamma}(t)_{i} - Z, t) - F(\hat{\gamma}(t)_{i}, t)\right] dG(Z). \end{aligned} \tag{5.1}$$

After the deadline has elapsed, the barrier resumes the form it takes in equation (3.9). At the instance when the deadline elapses, the approval barrier makes a discontinuous jump upwards, which we depict graphically in Figure 2. The deadline bonus results in a downward shift in the approval barrier, as the regulator requires a lower burden of proof induced by the payoff of $D$. Similarly, the discontinuous jump upward occurs as the regulator's willingness to approve a case returns to trajectory it would have without a deadline bonus.

Figure 2 About Here

Deadlines and deadline bonuses are usually imposed to decrease the time to approval in regulation. Our first set of comparative statics demonstrate that this is predicted by our formal model.

**Proposition 2** (Deadline Bonuses Decrease Time to Approval)**.** *For a set of cases $N, i \in N$ $i = 1, \ldots, n$, fix a set of regulatory histories $H_i(t)$. Suppose there are two deadline bonuses offered $D, D^{'}$, $0 \leq D^{'} < D$, and assume $t^{D} = t^{D^{'}}$. Call the (possibly infinite) approval time for case $i$ with bonus $D$, $t_{app}^{i,D}$. For all $i \in N$, $t_{app}^{i,D} \leq t_{app}^{i,D^{'}}$.*

Proposition 2 shows that the imposition of a deadline bonus will reduce the time to approval for a set of cases, and that larger deadline bonuses result in faster reviews, all things equal. The logic behind this result is straightforward and generalizes Proposition 3 of Carpenter (2004): as the bonus increases, a case will be approved earlier (if approved before the deadline) as the barrier shifts downward.

In addition to altering the size of the deadline bonus offered, politicians can modify the time at which the deadline bonus elapses. The next proposition demonstrates that, as the amount of time a deadline bonus is offered increases, the time to approval of cases is weakly decreasing.

**Proposition 3** (Deadline Durations Decrease Time to Approval)**.** *For a set of cases $N$, $i \in N$ $i = (1, \ldots, n)$, fix a set of regulatory histories $H_i(t)$. Suppose there are two potential times for deadlines to elapse $t^{D} > t^{D^{'}}$, and that the associated deadline bonuses are equal $D = D^{'}$. For all $i \in N$, $t_{app}^{i,D} \leq t_{app}^{i,D^{'}}$*

Like Proposition 2 this result is straightforward. As the deadline is extended, the approval time of a case that is approved before the deadline elapses remains unchanged. But, there are a set of regulatory histories that result in new cases being approved before the deadline elapses, as the deadline is pushed back. These cases will have a smaller approval time, proving the result.

Propositions 2 and 3 imply the following corollary, which shows that increasing the deadline bonus or extending the date a deadline bonus expires decreases the expected approval time for cases that are guaranteed to be approved.

**Corollary 1.** *Consider a case $i$ such that $\mu_i > \delta A$. Then the expected time to approval for case $i$, $E[t_{app}^{i,D}]$ is strictly decreasing in the deadline bonus and deadline time.*

# 6 The Trouble with Deadlines as a Tool of Control

The results of the previous section demonstrate that imposing deadlines upon an agency will reduce the expected time to approval for a case. This suggests that deadlines can be an effective tool of political control of regulation. But imposing deadlines upon an agency is complicated and results in unintended "side-effects" on the regulatory process. In this section we demonstrate that the relationship between the imposed deadline regime and regulatory behavior is highly complicated, limiting the ability of elected officials to effectively use deadlines to control agents. We demonstrate three specific results about deadlines: the expected approval time is non-linear in the deadline bonus and time, deadlines can make regulatory behavior less predictable, and deadlines increase the error rate of regulation. All three properties complicate deadlines as a tool of control over agencies.

## 6.1 Non-Linear

We use simulation to demonstrate the non-linear relationship between deadline bonuses and the expected time to approval. Figure 3 presents 100 simulations of a regulatory process and the corresponding time to approval, varying the deadline and the size of the deadline bonus. The left-hand figure is a contour plot of the expected time to approval. On the vertical axis the deadline bonus is varied and along the horizontal axis the deadline is varied. The colors represent the expected time to approval, with purple colors representing longer reviews of the drug and light-blue representing shorter reviews.

Figure 3 About Here

The left-hand plot in Figure 3 demonstrates the complicated relationship between the deadline, the size of the deadline bonus, and the expected time to approval. For example, if a deadline is offered for 26 units, then the expected time to approval can vary from 75.8 units (with deadline bonus 0.5) to an expected review time of 11.4 (with deadline bonus 10). Likewise, a deadline bonus of 6 can result in an expected review time of 63.9 periods when the bonus is offered for 10 periods and 35.4 periods when the bonus is offered for 30 periods.

Figure 3 also shows that the relationship between these points is highly nonlinear. To formally demonstrate this, we used the simulations from the left-hand plot to measure the *curvature* of the expected approval time in the time of deadline (center-plot) and the size of the deadline bonus (right-hand plot). Curvature allows for a formal measure of departures from linearity. For a twice-differentiable function $f : X \rightarrow Y$, The curvature of $f$ at $x \in X$, $c(x)$ is defined as (Shifrin, 2005),

$$c(x) \quad = \quad \frac{|f^{''}(x)|}{(1 + f^{'}(x))^{3/2}}$$

If the function is linear at $x$, then $c(x) = 0$ (the second derivative of any linear function is 0 everywhere) and as $c(x)$ increases, the function is increasingly non-linear at $x$. Both the center- and left-hand plot shows that the relationship between the expected approval time and the tools of control in a deadline are highly non-linear. The center plot shows that the curvature is greater than zero for almost all deadline times including in the simulation. Likewise, the curvature increases sharply as the size of the deadline bonuses increase.

The non-linear relationship between the expected approval time and the deadline bonus is problematic for elected officials who are attempting to use deadlines to control the behavior of agencies. Elected officials are unlikely to know exactly how a regulatory agency will respond to the imposition of a deadline bonus and deadline time. As a result, elected officials will have to estimate the likely effect of the deadline bonus. But officials will encounter substantial error if they attempt to linearly extrapolate from their information about an agent's reaction to the deadline bonus, lowering the effectiveness of deadline bonuses on controlling agency behavior. Likewise, this non-linear relationship may lead elected officials to conclude that deadlines are an ineffective tool of political control–when a small increase deadlines could allow a politician to substantially decrease the expected approval time.

## 6.2 Variance

Not only is it difficult for elected officials to assess the likely effect of a deadline bonus and deadline, under a wide range of conditions, the approval time will become more unpredictable after the imposition of a deadline. To demonstrate this result, we first prove Lemma 2 which will help us to characterize the approval time distribution.

**Lemma 2** (Multimodal Approval Distribution After Deadlines)**.** *Fix a deadline bonus $D, D > 0$ and a deadline time $t^D$. Call the approval time distribution for case $i$, $p(t^i_{app})$. Then there is $\varphi > 0$ such that $p(t) = 0$ for all $t \in [t^D, t^D + \varphi]$*

Figure 4 About Here

Lemma 2 demonstrates that imposing deadlines upon a regulator ensures that no cases are approved immediately after the deadline elapses, or that the approval distribution is (at least) bimodal. This suggests that we can model the effect of deadlines on the approval time distribution using a mixture of two densities: one distribution for the approval times of cases approved before the deadline and a second distribution for cases approved after the deadline. Formally, suppose that the approval time $t^i_{\text{app}}$ for a case $i$ is distributed according to the density $b(t^{i,D,t^D}_{\text{app}})$ if approved before the deadline bonus, with mean $\bar{t}^b_{D,t^D}$ and variance $\sigma^{2,b}_{D,t^D}$, where $D$ denotes the dependence on the size of the deadline bonus $D$ and $t^D$ the dependence on the actual deadline. If the case is approved after the deadline, suppose that its approval time is distributed $a(t^{i,D,t^D}_{\text{app}})$ with mean $\bar{t}^a_{D,t^D}$ and variance $\sigma^{2,a}_{D,t^D}$.[9] Further, suppose that case $i$ is approved before the deadline with probability $p^{i,D,t^D} \in (0,1)$. We can write the density of approval times $p(t^{i,D,t^D}_{\text{app}})$ as,

$$p(t^{i,D,t^D}_{\text{app}}) \quad = \quad p^{i,D,t^D} b(t^{i,D,t^D}_{\text{app}}) + (1 - p^{i,D,t^D}) a(t^{i,D,t^D}_{\text{app}}). \tag{6.1}$$

Equation 6.1 makes assessing the variance in the approval time distribution straightforward, using the law of total variation. Proposition 4 formalizes the conditions where increases in deadline times and bonuses will alter the variance in the approval distribution.

**Proposition 4** (Deadlines Can Render Regulation Less Predictable)**.** *Fix two deadline bonuses $D, D^{'}$ and deadlines $t^D, t^{D^{'}}$. Without loss of generality assume $D > D^{'}$ and $t^D > t^{D^{'}}$. Then the pair $(D, t^D)$ will*

[9] Assume both densities have support on $\Re_+$ and finite second moments.

*increase the variance in the approval time distribution of case $i$ over $(D^{'}, t^{D^{'}})$ if and only if*

$$
\begin{aligned}
0 \;<\;& p^{i,D,t^D}\sigma^{2,b}_{D,t^D} - p^{i,D^{'},t^{D^{'}}}\sigma^{2,b}_{D^{'},t^{D^{'}}} + \left(p^{i,D,t^D} - (p^{i,D,t^D})^2\right)(\bar{t}^{b}_{D,t^D})^2 - \left(p^{i,D^{'},t^{D^{'}}} - (p^{i,D^{'},t^{D^{'}}})^2\right)(\bar{t}^{b}_{D^{'},t^{D^{'}}})^2 \\
+\;& \left((1-p^{i,D,t^D}) - (1-p^{i,D,t^D})^2\right)(\bar{t}^{a}_{D,t^D})^2 - \left((1-p^{i,D^{'},t^{D^{'}}}) - (1-p^{i,D^{'},t^{D^{'}}})^2\right)(\bar{t}^{a}_{D^{'},t^{D^{'}}})^2 \\
+\;& (1-p^{i,D,t^D})\sigma^{2,a}_{D,t^D} - (1-p^{i,D^{'},t^{D^{'}}})\sigma^{2,a}_{D^{'},t^{D^{'}}} - 2p^{i,D,t^D}(1-p^{i,D,t^D})\bar{t}^{b}_{D,t^D}\bar{t}^{a}_{D,t^D} \\
+\;& 2p^{i,D^{'},t^{D^{'}}}(1-p^{i,D^{'},t^{D^{'}}})\bar{t}^{b}_{D^{'},t^{D^{'}}}\bar{t}^{a}_{D^{'},t^{D^{'}}} \qquad (6.2)
\end{aligned}
$$

Proposition 4 demonstrates that increasing the deadline bonus or extending a deadline can make administrative decision making less predictable. The proposition also demonstrates that assessing whether variance will increase after a change in the deadline bonus requires extensive knowledge about the regulatory process that an elected official is unlikely to possess while designing the deadline regime and would be difficult to estimate after the deadline is imposed. In short, Proposition 4 demonstrates that politicians are likely to have high uncertainty about whether a change in deadlines will make governmental decision making less predictable.

To demonstrate the complicated relationship between the parameters of the review process and the change in variance, Figure 6 evaluates Equation 6.2, applying different changes to the underlying mean, while allowing for a low-variance case for the component distributions (left-hand plot) and a high-variance case (right-hand plot). As the colors move from white to purple, the deadlines result in an increase in variance–and blue areas indicate that the combination of changes in means of the component distributions of the mixture decrease the variance in the approval distribution. Figure 6 shows that deadlines can increase the variance in the approval distribution, depending exactly on how the deadline affects the regulatory behavior. Furthermore, this figure shows that the effect of deadlines on the variance in approval times is complicated, making it unlikely elected officials would be able to intentionally alter the predictability of drug reviews through the imposition of deadlines.

Figure 6 About Here.

## 6.3 Error Rates

### 6.3.1 Defining Regulatory Error

The types of errors made in an inferential setting such as this one depend on the reference point or "null hypothesis" defined. If the null hypothesis is that the case should not be approved – that, for instance, the case submitted to the FDA is not "efficacious" until proven so – then a Type I error is the rejection of the null

by "approval" of the case when the proper decision should have been rejection or withholding of approval. A Type II error would then correspond to acceptance of the null hypothesis when it should have been rejected, that is, the rejection of, or failure to approve, a "good" case. If the null hypothesis is that the case should be approved, then this typology is inverted, with a Type I error corresponding to faulty rejection and a Type II error corresponding to faulty approval. In part due to convenience and consistency with earlier literature, and in part because many if not most regulatory procedures that involve optimal stopping are characterized by the first set-up – assume the case should not be approved until proven so, by some criteria – we adopt the null hypothesis that the case should not be approved until an evidentiary basis has been satisfied, and define Type I and Type II errors accordingly (see also Carpenter and Ting (2007), Bendor (1985), and Heimann (1993)).

The event that the case should or should not have been approved can be clearly described in terms of the model. The case should not be approved – according to the regulator's own objectives and goals – if the "true" value of the case lies below the value of rejection (or, equivalently, the value of infinite continuation). In terms of the parameters and variables of the present model, this corresponds to the event that $\mu_i(t) < \delta A$. Accordingly, the case should be approved if $\mu_i(t) > \delta A$. Note, that we do not consider the difference between the compensated jump process in the definition of error, because this is an expectation zero process. Any difference is due to randomness in the process and does not reflect fundamental characteristics of the case. [10]

Given the description of regulatory behavior that we have developed in the above model, we can then state the probabilities of Type I and Type II error as follows. Let $\Phi^I$ denote the probability of Type I error given the null hypothesis that the case should not be approved until proven effective, valid or otherwise "good." Define $\Phi^{II}$ as the corresponding probability of Type II error. Then

$$\Phi^I = \Pr\left[\exists\, t \in [0,\infty) \text{ s. t. } \hat{\gamma}(t)_i \geq \eta(t) | \mu_i < \delta A\right] \tag{6.3}$$

and

$$\Phi^{II} = \Pr\left[\nexists\, t \in [0,\infty) \text{ s. t. } \hat{\gamma}(t)_i > \eta(t) | \mu_i > \delta A\right] \tag{6.4}$$

[10] We will not deal with the case of equality here, as it is a knife-edge occurrence that has measure zero. It is sufficient to state a tie-breaking rule such that the case should be approved if $\mu(t)_i = \delta A$. As it turns out, the case of equality is a non-trivially difficult one to analyze in terms of dynamic stochastic movements.

It can be shown that the probability of Type I error is always non-zero in finite time, and given the current set-up of the model, there is no way of avoiding Type I error if the regulator's preferences and constraints allow it any discretion whatsoever. Similarly, Type II error rates are possible, at least if regulation is restricted to some finite, but large period of time, because the evidence process of a case may suffer a string of "unlucky" events, dooming a case that should be approved to endless consideration.

### 6.3.2 Deadlines Increase the Error Rate

Using this definition of regulatory error, we demonstrate that the probability of Type I error is weakly increasing in the deadline bonus.

**Proposition 5** (Deadlines Increase Error Rate)**.** *For any case $i$ such that $\mu_i < \delta A$, the probability of a type I error is weakly increasing in the deadline bonus $D$ and the total time the deadline bonus is offered, $t^D$.*

Proposition 5 presents a fundamental tension inherent in the application of deadlines: the tradeoff between decreased time-to-approval and increased error rates. Lowering the approval barrier decreases the time-to-approval for all drugs. However, lowering the approval barrier also results in more errors, because the set of cases approved are weakly increasing in the size of the deadline bonus, and therefore the set of bad cases that are approved is also weakly increasing in size as the deadline bonus increases.[11] This problem persists because elected officials lack the information to selectively apply deadline bonuses.

Given the tradeoff between decreased time-to-approval but increased error rates, politicians face a difficult institutional design problem: select the value of $D$ that optimizes the elected official's tradeoff between errors and faster regulation. Solving this problem is rendered more complicated because the actual error rate is highly non-linear in the size of the deadline bonus, and highly dependent upon the other parameters outlined in the model. The left-hand plot in Figure 5 presents simulations that demonstrate how the Type-I error rate depends upon the size of the deadline-bonus (the horizontal axis) and the expected jump-size (vertical axis), and the colors represent iso-probability curves for Type-I error. Notice, the curves are highly curved, indicating that increases in Type I error is highly non-linear in both jump-size and deadline bonus. This non-linearity in the deadline-bonus is also demonstrated in the right-hand plot in Figure 5, which shows the error rate as a function of the deadline bonus, my marginalizing the jump-size. Notice, that for low and high

[11] Fix a regulatory history. Define the set of drugs approved with deadline bonus $D$, $\mathcal{N}_D$. Now, suppose $D^{'} > D$. Then each drug $i \in \mathcal{N}_D$ must also be in $i \in \mathcal{N}_D^{'}$, otherwise we reach a contradiction that the barrier under $D$ is lower than $D^{'}$.

deadline bonuses, small increases in the deadline bonus result in essentially no increase in the type I error rate. But as bonuses increase from 11-13, the type-I error rates are highly responsive. Therefore, even if the politician could estimate the error rate after a deadline bonus, she might conclude that bonuses have little affect on error rates (say if the bonus is shifted at low values) and therefore leave the bonus unchanged.

Figure 5 About Here

# 7 Variance of Approval Time at the FDA

Our formal model makes the counterintuitive prediction that imposing a deadline bonus can *increase* the variance in approval times for government decision makers. To show that this is more than a theoretical possibility, we analyze the variance in approval times for new molecular entities submitted to the FDA and placed under priority review. The Prescription Drug User Fee Act (PDUFA) imposed deadlines on the decision-time during FDA reviews (Carpenter, Zucker and Avorn, 2008) and drugs under priority review had the most stringent deadlines. Therefore, approval times of priority drugs is a likely place to uncover the counterintuitive effects of deadlines on regulatory behavior.

Directly assessing how deadline regimes affected drug reviews under PDUFA is nearly impossible, because the PDUFA deadlines are conflated with a massive increase in FDA staff. To demonstrate the increase in staff size, the left-hand plot in Figure 7 shows that the FDA staff nearly doubled from 1988 to 2002. The vertical axis represents the mean-staff size, the horizontal axis varies the year, and the red-line is a non-parametric regression of the mean staff size on time. Clearly, each year the FDA increased the staff allocated to drug review substantially, which is conflated with the imposition of the deadline. Therefore, we look for suggestive evidence of deadlines' influence on the variance in approval times by using year trends in the variance of the approval time distribution.

Intuition and our theoretical analysis (Proposition 2 and Corollary 1), suggest an increase in staff size and imposing deadlines should decrease the mean time-to approval. The center-plot shows this relationship: the mean yearly approval time (vertical-axis) decreases substantially from 1988 to 2002. From 1988 to 2002, the mean approval time is lowered by over 5 months, a more than 25% reduction in mean approval time.

Intuition also suggests that the increase in the staff at the FDA should manifest in regulation becoming more predictable or a decrease in the variance of approval times. But, contrary to intuition and in line with our theoretical results (Proposition 4), the increases in staff are associated with *increased* variance in

approval times after the deadlines are imposed (right-hand plot). In this plot each point represents the variance in approval times for priority drugs submitted to the FDA in a given year, while the red-line represents a nonparametric regression of variance on time. This regression shows that regulatory decisions are *less predictable* after doubling staff for drug reviews and deadline reforms designed to better control regulatory decisions. While a number of causes for this unchanged variance in approval times are possible, we suggest that the most likely candidate are deadlines.

Figure 7 About Here

For evidence of deadline's impact on the approval time distribution, Figure 8 shows the approval time distribution for priority drugs, both before and after the imposition of the PDUFA deadlines. The left-hand plot shows that the time to approval distribution before the deadlines was relatively smooth and unimodal. By contrast the right-hand plot shows that deadlines create a bimodal distribution, which is predicted by Lemma 2. The sharp spike occurs as the deadline for priority reviews elapses, at 6 months. As argued above, this can induce an increase in the variance in the approval time distribution–which could negate the effects of increased staff on the variance in FDA approval times.

Figure 8 About Here

# 8 Empirical Implications and Conclusions

Our analysis highlights the necessity of careful consideration of how institutions affect regulatory behavior. Using a formal model of regulatory learning that is more general than has previously been considered, we show how deadline bonuses alter the regulator's behavior in the desired direction–decreasing the time to approval for all cases under review. Yet, we demonstrate that deadlines are not a simple tool of political control. First, the relationship between deadlines and the expected time to approval is highly non-linear, making the politician's decision about how to apply deadlines more complicated. Second, we show that deadlines can *increase* the level of uncertainty in the regulatory process. Third, we show the application of deadline bonuses to all cases under review induces increased error rates, as regulators are encouraged to offer faster review on cases that should not be approved.

In spite of the complicated relationship between deadlines and regulatory behavior, we have not shown that deadlines are net negative in terms of welfare, because the benefits from speeding up regulatory decisions may outweigh the costs of added error. The benefit can be quantified by thinking of the set of cases that

would eventually be approved (for which $\mu_i \geq \delta A$) but which are approved before the deadline and would otherwise have been approved afterwards. When this value exceeds the value lost from higher Type I error, then the deadline can be said to be welfare-improving within the constraints of the model. However, if there are other costs to deadlines (in that the benefits of the case are perhaps dependent upon the amount of time taken to learn about its parameters), or if there are other benefits to deadlines (in that the deadlines perhaps induce grater efficiencies by the regulator that spill over to other activities), then the policy calculation of the present model is inadequate and will fail to capture these benefits and costs.

### 8.1 Implications for Models and Tools of Political Control

Deadline institutions are commonly employed by politicians and courts as a way of constraining administrative and regulatory choice. Our model points to the effectiveness of these institutions, but also suggests that the effects of deadlines can be quite complicated. First, Proposition 4 shows that the imposition of a deadline upon an agency's decision-making process can actually render the agency's decision-making less predictable (or more variable), even as it reduces decision time for many of the agency's choices. In other words, far from making political control easier, deadlines may complicate political control of agency choice. Second, our analysis suggests that the imposition of new deadlines and deadline bonuses can generate greater administrative error. The relationship between agency payoffs (for meeting the deadline) and error rates is, moreover, highly dependent upon many unobservable and unpredictable variables (see Figure 5). It is difficult to believe that deadlines are (or can be) optimized *ex ante* to account for all of these contingencies.

We envision a productive research agenda at the intersection of deadline models and game-theoretic models of administrative politics, including delegation. Under what conditions would a principal wish to impose a deadline upon an agency decision-making process? How might an agency strategically react to the deadline institution? These are fruitful questions for further modeling.

### 8.2 Empirical Work on Deadlines and Administrative Error Rates

Our model also predicts a relationship between deadline-induced decisions and error rates. Carpenter, Zucker and Avorn (2008) examine deadlines imposed upon the FDA drug approval process and demonstrate that new drugs approved just before the deadlines are three to five times more likely to encounter measurable safety problems once on the market. Still, there is at present limited observational evidence pointing to a relationship between deadlines and error, and examination of data in varied settings is neces-

sary to more fully and accurately document this relationship. Moreover, Carpenter, Zucker and Avorn (2008) do not measure the deadline bonus (or payoff) associated with the FDA's having met a deadline. Nor do they measure other concepts embedded in the present model that may vary across time and across decisions, such as the population-level product danger ($\lambda_i$), the approval payoff ($A$) and the magnitude of the danger signals ($\xi$). One possibility is a regression discontinuity approach that compares the error rates for cases approved immediately before the deadline versus those that are approved immediately afterwards. If the difference between "just before" and "just after" is plausibly random, then a more experimental approach to assessing the relationship between deadlines and error is possible.

### 8.3 Empirical Implications for the Statistical Modeling of Bureaucratic Behavior

Our model also generates lessons concerning the empirical methods for studying the duration of administrative and regulatory decisions. Whereas learning processes characterized by continuous diffusions alone induce parametric forms for the distributions governing regulatory decision time (e.g., the inverse Gaussian, see Carpenter (2002, 2004)), our model predicts a hazard function which violates both standard parametric assumptions for a hazard (e.g. exponential, gamma, inverse gaussian) and the proportional hazards assumptions often used in semi-parametric analysis of duration data (for example, the Cox model). To see why, we first prove in Corollary 2 that deadlines cause a discontinuity in the hazard rate.

**Corollary 2.** *For a case $i$, the hazard function $h_i(t)$ drops to zero immediately after the deadline elapses, for all $D > 0$. Formally,* $\lim_{t \downarrow t^D} h_i(t|t > t^D) = 0$

The violation of the proportional hazards assumption occurs because the hazard rate drops to zero as the deadline elapses, and therefore the coefficients on all covariates must drop to zero as well at this instance. We conjecture that shifts in behavior after a deadline elapses may further exacerbate the violation of proportional hazards. We leave analysis of a more appropriate estimator to another paper. For now, suffice it to say that the introduction of greater realism into the evidence processes about which regulators learn yields much greater complexity of behavior than is appreciated or embedded in standard statistical models.

### 8.4 Extensions and Future Work

Two limitations of this modeling framework strike us as ideal targets for thoughtful extension. First, a primary determinant of time-to-decision in organizational settings is not simply the duration of "optimal stopping," but organizational features such as the queue of cases coming to the agency. These flows may

depend upon strategic considerations (Carpenter and Ting, 2007), while in other cases factors such as organizational efficiency and the number of organizational units reviewing cases may also be influential (Bendor, 1985; Heimann, 1993; Ting, 2003). Such queueing processes have been well studied in stochastic analysis, but to our awareness, models which embed *both* stopping behavior and queues have not been attempted. Second, it is quite possible that much of administrative decision making amounts not simply to optimal stopping but also to "optimal control," in that the eventual quality or danger of the case may in fact depend upon the amount of time that the regulator or risk analyst has spent learning about it. Whenever this is true, the signal extraction metaphor governing our model leaves much to be desired. We suspect that introduction of queues and case-based optimization will yield rich theoretical progress, and until this happens, the conclusions of the present analysis should be taken with circumspection.

# A Appendix

The regulator observes the unfolding of evidence on a space $\Omega$ (with elements or experimental realizations $\omega$), which is structured by a set of $\sigma$-algebras $\Im$, and a probability measure $P$. In addition, $\Im$ can be ordered and expressed as a filtration $(\Im_t)_{0\leq t\leq\infty}$, which is a family of $\sigma$-algebras that is increasing in its index, hence $\Im_s \subset \Im_t$ if $s \leq t$. The filtration sequentially collects and orders all realizations $\omega = \omega_t$ on a time dimension from 0 to $t$. The collection $(\Omega, \Im, \Im_t, P)$ constitutes a filtered probability space, on which we assume that a set of "usual hypotheses" standard in the analysis of stochastic differential equations. These hypotheses and a relatively clear explanation of their importance appear in Protter (2005: Chapter I, esp. pp. 34-36).

Using this more formal structure, we can make a more rigorous statement of Lemma 1.

**Lemma 1** *Let $\mathbf{F}_t = (\Im_t)$ represent the filtration for the evidence process $X(t)$ as given in Equation 3.3. Without loss of generality, for any $X(t)_{t\geq 0}$, any $\Im_t$ can be broken into three separable and independent components: (1) the filtration of the continuous diffusion, $\mathbf{F}_t^B$, (2)the filtration of the compensatory $\mathbf{F}_t^c$, and (3) the filtration of the jumps $\mathbf{F}_t^J$. Then a sufficient statistic for $\mathbf{F}_t^B$ is the dual $(t, X(t)^*)$, where $X(t)^* = X(t) - \lambda\xi t + \sum_{k=1}^{J(t)} Z_k$, a sufficient statistic for $\mathbf{F}_t^J$ is $(t, \sum_{k=1}^{J_t} Z_k > 0)$ and a sufficient statistic for $\mathbf{F}^c$ is (t, $\lambda_{i,j}\xi$)*

*Proof of Lemma 1.* Define $\Delta X(t) \equiv X(t) - X(t-)$ where $X(t-) \equiv \lim_{s\to t, s\leq t} X(s)$ and let $\mathcal{T}$ denote the set of stopping times. The agent knows the rate of compensation for the jump process, so we can "detrend" the evidence process to obtain $X^*(t) = X(t) - \lambda\xi t$ and knowledge of the compensation rate is obviously

sufficient to learn about $\mathbf{F}^c$. During the continuous portion of the process $\Delta X^*(t) = 0$ and during the instant where a jump occurs $|\Delta X^*(t)| > 0$. Define $\Delta X^{j,*}(t) \equiv \{t \in \mathcal{T} | \, |\Delta X^*(t)| > 0\}$- the set of jump times-and define $\Delta X^{B,*}(t) \equiv \{t \in \mathcal{T} | \, \Delta X^*(t) = 0\}$- the non-jump periods. These two sets are disjoint, so that information about each component arrives at separate times. Now, take $\mathbf{F}_t^B \equiv \{\mathbf{F}_s | s \in \Delta X^{B,*}(t)\}$ and $\mathbf{F}_t^J \equiv \{\mathbf{F}_s | s \in \Delta X^{J,*}(t)\}$. By assumption the jump process and the drift component is independent, which implies that $\mathbf{F}^B \perp \mathbf{F}^J$. Because the compensation rate is known to the agent, $\mathbf{F}^B, \mathbf{F}^J \perp \mathbf{F}^c$.

Note that $\mathbf{F}^B$ and $\mathbf{F}^J$ contain all the relevant information about the evidence process at time $t$ and because all Levy-Processes are Markov chains this is all the information necessary to understand how the evidence history predicts the future of the process. It follows directly that $(t, X(t))$ is a sufficient statistic for $\mathbf{F}^B$ and that $(t, \sum_{k=1}^{J_t} Z_k)$ is sufficient to summarize the information in $\mathbf{F}_t^J$. □

*Proof of Proposition 1.* The proof proceeds in three parts. First, we prove that the filtered evidence process $\hat{\gamma}(t)_i$ is a Levy process. Then, we rewrite the problem in order to remove the components of the expression that will equal zero at the barrier. Finally, we utilize smooth pasting and value matching conditions to determine the optimal stopping barrier.

We first show that the filtered evidence process $\hat{\gamma}(t)_i$ is a Levy process.

**Lemma 3.** *$\hat{\gamma}(t)_i$ is a Levy Process.*

*Proof.* As Miroschnichenko (1975) shows, $\hat{\mu}(t)_i$ is rescaled Brownian motion. Therefore, $\hat{\mu}_i(t)$ is itself a Levy process (Protter, 2005; Applebaum, 2004). Now, as in Thm I. 37 in Protter (2005), we need only check for stationary increments of our process, $\hat{\gamma}_i(t)$. Expressing the difference,

$$\begin{aligned} \hat{\gamma}(t)_i - \hat{\gamma}(s)_i &= \hat{\mu}(t)_i - \sum_{k=0}^{J_t} Z_k + \lambda_i \xi t - (\hat{\mu}(s)_i - \sum_{k=0}^{J_s} Z_k + \lambda_i \xi s) \\ &= \hat{\mu}(t)_i - \hat{\mu}(s)_i + \lambda \xi (t-s) + \sum_{s<u\leq t} \Delta \hat{\gamma}(u)_i. \end{aligned} \tag{A.1}$$

Equation (A.1) is $\sigma\{X(v) - X(u) : s \leq u < v \leq t\}$ measurable, because all three components of Equation (A.1) are derived directly from $X(t)$. Therefore, $X(t)$ has independent increments, it follows that $\hat{\gamma}(t)_i$ has independent increments and is therefore also a Levy process.[12] □

[12]Some texts use the definition of a Levy process as the additive combination of a Brownian motion (with drift) and a compensated Poisson process. In which case, the claim of Lemma 2 would follow by definition.

Writing the expression for the derivative of $F(\hat{\gamma}(t)_i, t)$,

$$\delta F(\hat{\gamma}(t)_i, t) \quad = E[df] \quad = E_{\hat{\gamma}(t)_i, t}\{F(\hat{\gamma}(t)_i, t + dt) - F(\hat{\gamma}(t)_i, t)\} + o(dt) \tag{A.2}$$

where $o(dt)$ collects terms that vanish as $dt$ goes to zero. We use a fundamental result from stochastic calculus, the Meyer-Ito formula, an analog of Taylor-series expansion to rewrite Equation (A.2). Using this result, we will be able to express Equation (A.2) in a form familiar to analysts, $\delta F(\hat{\gamma}(t)_i, t) =$ Process + Error.

To apply the Meyer-Ito Lemma, $F(\cdot, \cdot)$ must be twice continuously differentiable, which we have assumed and $\hat{\gamma}(t)_i$ must be a semimartingale, which follows from Lemma 2 and Thm II.9 and its corollary in Protter 2005, 55-56.[13] Rewriting Equation (A.2), we have,

$$\begin{aligned} F(\hat{\gamma}(t)_i) - F(\hat{\gamma}(0)_i) \quad = \quad & \int_0^t F^{'}(\hat{\gamma}(s-)_i) d\hat{\gamma}(s)_i + \frac{1}{2} \int_0^t F^{''}(\hat{\gamma}(s-)_i) d \left\langle \hat{\gamma}(s)_i, \hat{\gamma}(s)_i \right\rangle_s^c \\ & + \sum_{0<s\leq t} \{F(\hat{\gamma}(s)_i) - F(\hat{\gamma}(s-)_i) - F^{'}(\hat{\gamma}(s-)_i) \Delta\hat{\gamma}(s)_i\}, \end{aligned} \tag{A.3}$$

where $\langle \cdot, \cdot \rangle_s^c$ of the continuous portion of the $\hat{\gamma}(s)_i$, which will be equal to the variance of $\hat{\mu}_i$ (Protter, 2005; Applebaum, 2004).

The next step is to rewrite the first order terms of in equation (A.3) (Mordecki, 1999). To do so, we first define two measures corresponding to the jumping process and its intensity. Define $\kappa = \kappa(\omega, d\hat{\gamma}, ds)$ as the jump measure for $X(t)$ (a measure that describes the size of the jumps) and define $\nu = \nu(d\hat{\gamma}, s) = \lambda ds G(d\hat{\gamma})$ its compensator. Using these two measures, we are able to rewrite the the first order terms of equation (A.3) as (Applebaum (2004) Thm 4.4.7, pg 226, Thm 4.4.10, pg 229),

$$\int_0^t F'(\hat{\gamma}(s-)_i) d\hat{\gamma}(s)_i + \sum_{0\leq s\leq t} \left(F(\hat{\gamma}(s)_i) - F(\hat{\gamma}(s-)_i) - F'(\hat{\gamma}(s-)_i) \Delta\hat{\gamma}(s)_i\right) =$$

$$\underbrace{\int_0^t F'(\hat{\gamma}(s-)_i) d\hat{\gamma}(s)_i^c}_{\text{Continuous Evidence}} + \underbrace{\int_0^t \int_{\Re} \{F(\hat{\gamma}(s-)_i - x) - F(\hat{\gamma}(s-)_i)\} \times (\kappa(\omega, dx \times ds) - \nu(dx \times ds))}_{\text{Jump Process (martingale)}}$$

$$+ \underbrace{\int_0^t \int_{\Re} \{F(\hat{\gamma}(s-)_i - x) - F(\hat{\gamma}(s-)_i)\} \times \nu(dx \times ds))}_{\text{Jump Process (systematic)}}.$$

[13] For notational clarity, we suppress the dependence upon $t$ in the following derivations.

The first integral allows us to rewrite the first order continuous portion, while the two double-integrals correspond to the jump portion of the process. We are now in a position to rewrite the difference of our value function using the following three terms,

$$F(\hat{\gamma}(t)_i) - F(\hat{\gamma}(0)_i) = \int_0^t L^{\hat{\gamma}} F(\hat{\gamma}(s)_i) ds + M(\hat{\gamma}(t)_i)_t + o(dt) \tag{A.4}$$

where $L^{\hat{\gamma}} F(\hat{\gamma}(s)_i)$ is the infinitesimal generator of $\hat{\gamma}(s)_i$ (which is the continuous time analog of a Markov transition matrix), $M(s)_t$ is a local martingale and o(dt) collects the terms that go to zero as $dt \to 0$. By the above argument, we can express the $M(s)_t$ terms as,

$$M(\hat{\gamma}(s)_i)_t \quad = \quad \int_0^t F'(\hat{\gamma}) d\hat{\mu}_s + \int_0^t \int_{\Re} F(\hat{\gamma}_{s-} - x) - F(\hat{\gamma}_{s-}) \times (\kappa - \nu)$$

where we have suppressed the inputs of $\kappa$ and $\nu$ for clarity. Note, that the $\hat{\mu}_i(t)$ is a martingale because it is a (rescaled) Brownian motion and $\int_0^t \int_{\Re} F(\hat{\gamma}_{s-} - x) - F(\hat{\gamma}_{s-}) \times (\kappa - \nu)$ is a local martingale by the Doob-Meyer Decomposition Theorem.

Once again, from the argument above we can express the functional form for $L^{\hat{\gamma}} F(\hat{\gamma}(t)_i)$ as (Applebaum (2004) pg 141, Kou and Wang (2003) pg 507),

$$\begin{aligned} L^{\hat{\gamma}_t}(F)(\hat{\gamma}, t) \quad &= \quad F_{\hat{\mu}}(\hat{\gamma}(t)_i, t) + F_t(\hat{\gamma}(t)_i, t) + \frac{1}{2} V(t) F_{\hat{\mu},\hat{\mu}}(\hat{\gamma}(t)_i, t) \\ &\quad + \lambda \int_{\Re} F(\hat{\gamma}(t)_i - Z, t) - F(\hat{\gamma}(t)_i, t) dG(Z). \end{aligned} \tag{A.5}$$

Lemma 1 of Mordecki (1999) shows a localizing sequence can be defined showing that the $M(\hat{\gamma}(t^*)) = 0$ at the optimal stopping time, $t^*$.[14] Furthermore, as $dt \to 0$, $o(dt)$ vanishes. Therefore, we can now focus our attention exclusively on the infinitesimal generator in equation (A.5).

Upon restricting our attention to the infinitesimal generator, we note that, as in Kyprianou and Suryi (2005) and Novikov and Shiryaev (2004) that discrete jumps only occur downward and therefore cannot cause the process to cross an optimal barrier.[15] Therefore, a combination of value matching and smooth pasting conditions will be sufficient to characterize the free boundary.

For the value matching condition we have,

$$\delta F(\hat{\gamma}(t)_i, t) \quad = \quad \delta A + \hat{\gamma}^*(t)_i \tag{A.6}$$

[14] A localizing sequence is a set of stopping times such that some property of a stochastic process holds. This is analogous to the notion of locality in Real analysis where a process is restricted to some interval on the real line.

[15] With a more general setting where jumps occur in both directions, we would need to focus on the problem of "overshoot" or allowing an actor to account for the fact that the process may jump over a barrier rather than reach the barrier by the continuous information from the drift and Brownian motion. See Kou and Wang (2003) for a particularly clear discussion and analysis of overshoot.

where $\gamma^*(t)_i$ denotes the values of the evidence process at the time of optimal stopping. For the smooth pasting condition, we have (Dixit and Pindyck, 1994, 210)

$$F_t(\hat{\gamma}(t)_i, t) \quad = \quad \delta\lambda\xi \tag{A.7}$$

Note, that because $\hat{\mu}_i(t)$ is a martingale, $F_{\hat{\mu}}$ will also be a mean zero martingale (Miroschnichenko, 1975; Carpenter, 2002, 2004). Remembering that the local martingale is at zero at the point of optimal stopping, we rewrite Equation (A.2),

$$\begin{aligned} F_{\hat{\mu}}(\hat{\gamma}(t)_i, t) + F_t(\hat{\gamma}(t)_i, t) + \frac{1}{2\sigma^2}V(t)^2 F_{\hat{\mu},\hat{\mu}} - \delta F(\hat{\gamma}(t)_i, t) \\ +\lambda \int_{\Re} F(\hat{\gamma}(t)_i - Z, t) - F(\hat{\gamma}(t)_i, t) dG(Z) = 0. \end{aligned} \tag{A.8}$$

Substituting Equations (A.6) and (A.7) into Equation (A.8), yields,

$$\eta^* = \delta(\lambda\xi - A) + \frac{1}{2\sigma^2}V(t)^2 F_{\hat{\mu},\hat{\mu}}(\hat{\gamma}(t)_i, t) + \lambda \int_{\Re_+} \left[F(\hat{\gamma}(t)_i - Z, t) - F(\hat{\gamma}(t)_i, t)\right] dG(Z), \tag{A.9}$$

and this is what was desired to be shown, therefore we have completed the proof. □

*Proof of Proposition 2.* Fix $D, D^{'}$ such that $0 \leq D^{'} < D$ and a set of regulatory histories for all $i \in N$, $\Im_{i,t}$. We proceed by contradiction. Suppose that there is some case $i$ such that $t^{i,D}_{\text{app}} > t^{i,D^{'}}_{\text{app}}$. This implies that, at time $t^{i,D}_{\text{app}}$,

$$\eta^{*,D^{'}} \leq \hat{\gamma}(t^{i,D}_{\text{app}})_i \leq \eta^{*,D}$$

or that $\eta^{*,D^{'}} \leq \eta^{*,D}$. Because we have fixed the regulatory histories this implies that $\delta(D^{'} - D) \geq 0$, which contradicts our assumption that $0 \leq D^{'} < D$. We conclude that $t^{i,D}_{\text{app}} \geq t^{i,D^{'}}_{\text{app}}$. □

*Proof of Proposition 3.* Fix $t^D, t^{D^{'}}$ such that $0 \leq t^{D^{'}} < t^D$, assume $D = D^{'}$ and fix a set of regulatory histories $\Im_{i,t}$. Once again, we proceed by contradiction. Suppose there is $i$ such that $t^{i,D}_{\text{app}} > t^{i,D^{'}}_{\text{app}}$, which implies that $\eta^{*,D^{'}}_t \leq \hat{\gamma}(t^{i,D}_{\text{app}})_i < \eta^{*,D}_t$. But this cannot be before $t^{D^{'}}$, because $\eta^{*,D^{'}}_t = \eta^{*,D}_t$ for all $t \in [0, t^{D^{'}})$. And after $t^{D^{'}}$, $\eta^{*,D}_t > \eta^{*,D^{'}}_t$. Therefore, we have arrived at a contradiction and conclude that $t^{i,D}_{\text{app}} \geq t^{i,D^{'}}_{\text{app}}$. □

*Proof of Corollary 1.* Fix case $i$ fix a set of regulatory histories and assume $\mu_i > \delta A$. We prove the result for the deadline bonuses, the deadline times follow analogously. Assume $D > D^{'}$ By iterated expectations,

$\mathrm{E}[t^{i,D}_{\text{app}}] = \mathrm{E}[\mathrm{E}[t^{i,D}_{\text{app}}|\Im]]$ and that $\mathrm{E}[t^{i,D'}_{\text{app}}] = \mathrm{E}[\mathrm{E}[t^{i,D'}_{\text{app}}|\Im]]$. From Proposition 2 we know that $\mathrm{E}[t^{i,D}_{\text{app}}|\Im] \leq \mathrm{E}[t^{i,D'}_{\text{app}}|\Im]$ with the inequality strict for some $\Im$. Therefore, $\mathrm{E}[\mathrm{E}[t^{i,D}_{\text{app}}|\Im]] < \mathrm{E}[\mathrm{E}[t^{i,D'}_{\text{app}}|\Im]]$. □

*Proof of Lemma 2.* Without loss of generality, consider an evidence process immediately before the deadline elapses $\hat{\gamma}(t^{D-})_i$ that is arbitrarily close to the stopping barrier $\hat{\gamma}(t^{D-})_i + \zeta = \eta(t^{D-})$, $\zeta > 0$, that has yet to be approved. After the deadline elapses the barrier makes a discontinuous jump upward of size $\delta D$. Because the evidence process is continuous except for *downward* jumps, there exists $\epsilon, \varphi > 0$ such that $|t^{D-} - (t^{D-} + \varphi)| < \epsilon$ implies that $\Pr(|\hat{\gamma}(t^D_-) - \hat{\gamma}(t^D_- + \varphi)| < \delta D + \zeta) = 1$ or that, with probability 1, the evidence process will not cross the barrier for any $t \in [t^{D-}, t^{D-} + \varphi]$. □

*Proof of Proposition 4.* The proof proceeds by applying the law of total variability. We will denote whether a case is approved before or after the deadline with the random variable $Y$. For deadline bonus $D$ and deadline $t^D$ the variability is,

$$\text{var}(t^{i,D,t^D}_{\text{app}}) \quad = \quad \mathrm{E}\left[\text{var}(t^{i,D,t^D}_{\text{app}}|Y)\right] + \text{var}\left[\mathrm{E}[t^{i,D,t^D}_{\text{app}}|Y]\right]. \tag{A.10}$$

Carrying out the calculations for Equation A.10 shows that this is equal to

$$\begin{aligned}\text{var}(t^{i,D,t^D}_{\text{app}}) \quad &= \quad p^{i,D,t^D}\sigma^{2,b}_{D,t^D} + (p^{i,D,t^D} - (p^{i,D,t^D})^2)(\bar{t}^b_{D,t^D})^2 \\ &+ \quad \left((1-p^{i,D,t^D}) - (1-p^{i,D,t^D})^2\right)(\bar{t}^a_{D,t^D})^2 + (1-p^{i,D,t^D})\sigma^{2,a}_{D,t^D} - 2p^{i,D,t^D}(1-p^{i,D,t^D})\bar{t}^b_{D,t^D}\bar{t}^a_{D,t^D}.\end{aligned}$$

Consider deadline bonus $D'$ and deadline $t^D$. Then, the variance increases if $\text{var}(t^{i,D,t^D}_{\text{app}}) - \text{var}(t^{i,D',t^{D'}}_{\text{app}}) > 0$, which completes the proof. □

*Proof of Proposition 5.* We prove the result for increasing deadline bonuses $D$. The same argument holds for extensions of the deadline $t^D$.

Consider a case $i$ such that $\mu_i < \delta A$ and a review history $\Im_t$. Consider two barriers $\eta^*_{D,t^D}(t)$ and $\eta^*_{D',t^{D'}}(t)$, where the dependence on the deadline approval barrier is expressed in the subscript. Suppose, without loss of generality, $0 \leq D < D'$. Then, this implies that $\eta^*_{D'}(t) < \eta^*_D(t)$ for all $t$, as the barrier is decreasing in the size of the deadline bonus and we have fixed the case and history across the borders. Now, call the event of a type I error (the case approved) $\theta^D$. $\mathrm{E}[\theta^{D'}|\Im_t] > \mathrm{E}[\theta^D|\Im_t]$ for all $\Im_t$, because the barriers are decreasing in the size of $D$. Then, by iterated expectations,

$$\begin{aligned}\Phi^I_{D'} = E[\theta^{D'}] \quad &= \quad E[E[\theta^{D'}|\Im_t]] \\ &\geq \quad E[E[\theta^D|\Im_t]] = \Phi^I_D\end{aligned}$$

or that the error rate is weakly increasing as the size of the deadline bonus increases. □

*Proof of Corollary 2.* To prove Corollary 2, we first reexpress the approval problem as a counting process. Then, relying upon the Doob-Meyer decomposition theorem makes the conclusions of the theorem immediate.

Define $N(t)$ as a counting process adapted to $\Im_t$, with

$$N(t) \quad = \quad \begin{cases} 1 & \text{if } \hat{\gamma}(t^{'})_i \geq \eta(t^{'}) \text{ for some } t^{'} < t \\ 0 & \text{if } \hat{\gamma}(t^{''})_i < \eta(t^{''}) \text{ for all } \ t^{''} < t \end{cases} \quad . \tag{A.11}$$

$N(t)$ is a counting process because it starts at zero, jumps only once (and is therefore piecewise continuous) and is adapted to a right-continuous filtration by use of the usual assumptions, which we have assumed about $X(t)$.

We can apply the Doob-Meyer theorem to decompose $N(t)$ such that,

$$N(t) \quad = \quad \int_0^t h(t)dt + M_t \tag{A.12}$$

where $\int_0^t h(t)dt$ is the cumulative hazard rate and $M_t$ is a martingale. We can now define the hazard rate for case $i$ as

$$\mathrm{E}\left[dN_i(t)|\Im_{t-}\right] \quad = \quad r_i(t)h(t)dt \tag{A.13}$$

where $r_i(t)$ is a risk indicator which is equal to 1 if the case has not been approved and 0 otherwise. Because $N_i(t)$ is a counting process, $dN_i(t)$ is equal to either 0 or 1 and therefore we can restate the hazard rate as $h(t)dt = \Pr(dN_{i,j}(t) = 1|\Im_{t-})$ (Therneau and Grambsch, 2000). Therefore, we can prove Proposition 2 by analyzing the probability that the case crosses the barrier at a given instance, given that it has not yet crossed the barrier.

To complete the proof, without loss of generality, consider an evidence process immediately before the deadline elapses $\hat{\gamma}(t^{D-})_i$ that is arbitrarily close to the stopping barrier $\hat{\gamma}(t^{D-})_i + \zeta = \eta(t^{D-})$, $\zeta > 0$, that has yet to be approved. After the deadline elapses the barrier makes a discontinuous jump upward of size $\delta D$. Because the evidence process is continuous except for *downward* jumps, there exists $b, k > 0$ such that $|t^{D-} - (t^{D-} + k)| < b$ implies that $\Pr(|\hat{\gamma}(t_-^D) - \hat{\gamma}(t_-^D + k)| < \delta D + \zeta) = 1$ or that, with probability 1, the evidence process will not cross the barrier for any $t \in [t^{D-}, t^{D-} + k]$ and therefore $h(t)dt = 0$. □

## References


Abbott, Alden. 1987. "Case Studies on the Costs of Federal Statutory and Judicial Deadlines." *Administrative Law Review* 39:467.

Applebaum, David. 2004. *Levy Processes and Stochastic Calculus*. Cambridge University Press.

Bendor, Jonathan. 1985. *Parallel Systems: Redundancy in Government*. University of California Press.

Berns, Gregory, David Laibson and George Loewenstein. 2007. "Intertemporal Choice–Toward an Integrative Framework." *Trends in Cognitive Sciences* 11(11):482–488.

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

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

Carpenter, Daniel, Evan James Zucker and Jerry Avorn. 2008. "Drug-Review Deadlines and Safety Problems." *New England Journal of Medicine* 358(13):1354.

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

Dixit, Avinash and Robert Pindyck. 1994. *Investment Under Uncertainty*. Princeton University Press.

Fershtman, Chaim and Daniel Seidmann. 1993. "Deadline Effects and Inefficient Delay in Bargaining with Endogenous Commitment." *Journal of Economic Theory* 60:306–321.

Gersen, Jacob and Anne Joseph O'Connell. 2008. "Deadlines in Administrative Law." *University of Pennsylvania Law Review* Forthcoming.

Gneezy, Uri et al. 2003. "Bargaining Under a Deadline: Evidence from the Reverse Ultimatum Game." *Games and Economic Behavior* 45:347.

Gordon, Sanford and Catherine Hafer. 2005. "Flexing Muscle: Corporate Political Expenditure as Signals to Bureaucracy." *American Political Science Review* 99(2).

Heimann, Larry. 1993. "Understanding the Challenger Disaster: Organizational Structure and the Design of Reliable Systems." *American Political Science Review* 87(2):421–435.

Kosnik, Lea-Rachel D. 2006. "Sources of Bureaucratic Delay: A Case Study of FERC Dam Relicensing." *The Journal of Law, Economics, and Organization* 22(1):258–288.

Kou, Sam and Hui Wang. 2003. "First Passage Times of A Jump Diffusion Process." *Advances in Applied Probability* 35:504–531.

Kyprianou, Andreas and Budhi Suryi. 2005. "On the Novikov-Shiryaev Optimal Stopping Problems in Continuous Time." *Electronic Communications in Probability* 10:146–154.

Mashaw, Jerry and David Harfst. 1990. *The Struggle for Auto Safety*. Harvard University Press.

Miroschnichenko, TP. 1975. "Optimal Stopping of an Integral of a Wiener Process." *Theory of Probability and Its Applications* .

Moore, Don. 2004. "The Unexpected Benefits of Final Deadlines in Negotiation." *Journal of Experimental Social Psychology* 40:121.

Mordecki, Ernesto. 1999. "Optimal Stopping for a Diffusion with Jumps." *Finance and Stochastics* 3:227–236.

Morgenstern, Richard. 1993. Science, Engineering, and Regulation. In *Keeping Pace with Science and Engineering*. National Academic Press pp. 243–250.

Novikov, A and A Shiryaev. 2004. "On an Effective Solution fo the Optimal Stopping Problem for Random Walks." *Theory of Probability and Its Applications* 48.

O'Connell, Anne Joseph. 2008. "Political Cycles of Rulemaking: An Empirical Portrait of the Modern Administrative State." *Virginia Law Review* Forthcoming.

Olson, Mary. 1997. "Firm Characteristics and the Speed of FDA Approval." *Journal of Economics and Management Strategy* 6:377–401.

Protter, Phillip E. 2005. *Stochastic Integration and Differential Equations*. Springer.

Shifrin, Theodore. 2005. *Multivariable Mathematics: Linear Algebra, Multivariable Calcululs, and Manifolds*. John Wiley and Sons.

Spence, Michael. 1999. "Managing Delegation Ex Ante: Using Law to Steer Administrative Agencies." *Journal of Legal Studies* 28:413–459.

Stuhlmacher, Alice and Matthew Champagne. 2000. "The Impact of Time Pressure and Information on Negotiation Process and Decision." *Group Decision and Negotiation* 9:471.

Therneau, Terry and Patricia Grambsch. 2000. *Modeling Survival Data: Extending the Cox Model*. Springer.

Ting, Michael. 2003. "A Strategic Theory of Bureaucratic Redundancy." *American Journal of Political Science* 47(2):274–292.

Whitford, Andrew. 2005. "The Pursuit of Political Control by Multiple Principals." *Journal of Politics* 67(1):28–49.

Yackee, Jason and Susan Yackee. 2008. "Is Federal Rulemaking "Ossified"? The Effects of Congressional, Presidential, and Judicial Oversight on the Agency Policymaking Process.".

Figure 1: Visualizing The Optimal Stopping Problem

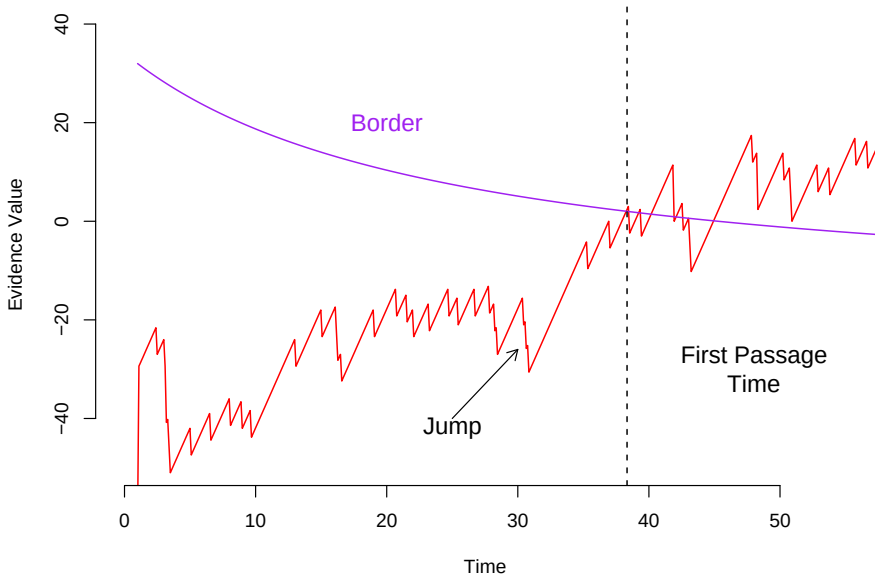


This figure visualizes the filtered evidence process and the optimal stopping barrier from Proposition 1. The horizontal axis represents time, the vertical axis is the utility provided to the regulator, the red-line is the filtered evidence process for one product review, and the purple line represents the optimal stopping barrier. Adverse events appear as a discontinuous jumps downwards, reflecting the regulators' aversion to approving cases when there is an indication that a drug may have an exceptionally poor safety record for public use. The border slopes downward, as the value of more information decreases over the course of the regulatory history. A case is approved only if its evidence crosses the boundary, which occurs at the right-hand side of the figure.

Figure 2: Deadlines Cause a Discontinuous Jump in the Approval Barrier

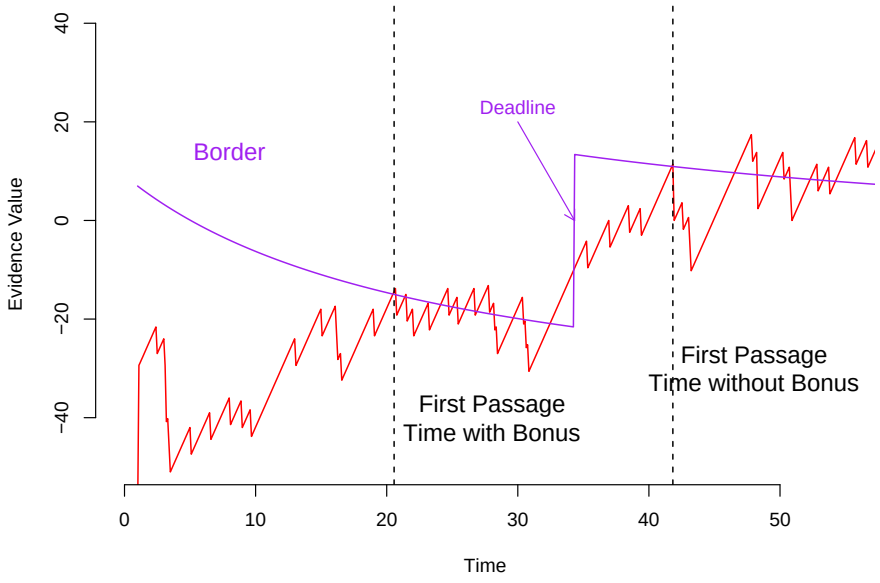


This figure presents an evidence process and barrier with a deadline imposed upon the regulator. The horizontal axis is time, the vertical axis the utility to the regulator for approval, the redline is the filtered evidence process and the purple line represents the optimal stopping barrier. Notice the large, discontinuous jump in this barrier after a deadline for case review elapses. This simple extension of the model causes a drastic change in regulatory behavior, which we detail below.

Figure 3: Extending the Deadline and Increasing the Deadline Bonus has a Non-Linear Effect on the Expected Approval Time

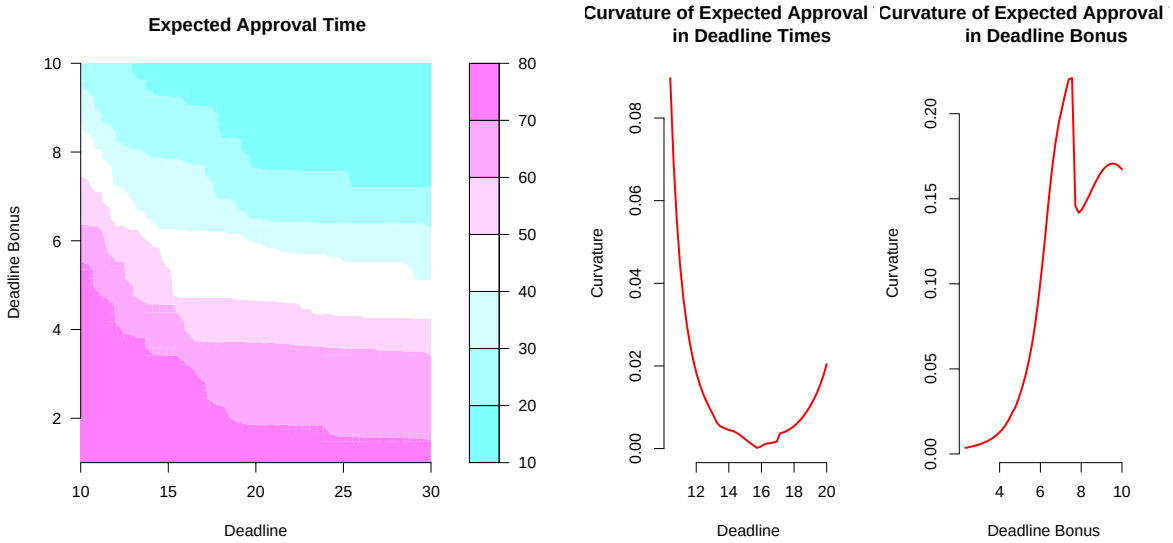


This Figure demonstrates the non-linear relationship between deadline bonuses and approval times and the expected approval time for a case under regulation. The left-hand plot shows how the expected approval time varies in a non-linear fashion and is dependent upon both the timing and size of the deadline bonus. To explicitly measure the non-linearity, the center- and right-hand plots measure the *curvature* of the expected approval time in the time of the deadline (center plot) and the size of the bonus (right-hand plot). If there is no curvature–or there were a linear relationship–then the curves in each figure would be at zero. But, we see here that the lines are highly curved, indicating a highly non-linear relationship.

Figure 4: Deadlines Cause Bimodal Distributions and Can Increase Variance

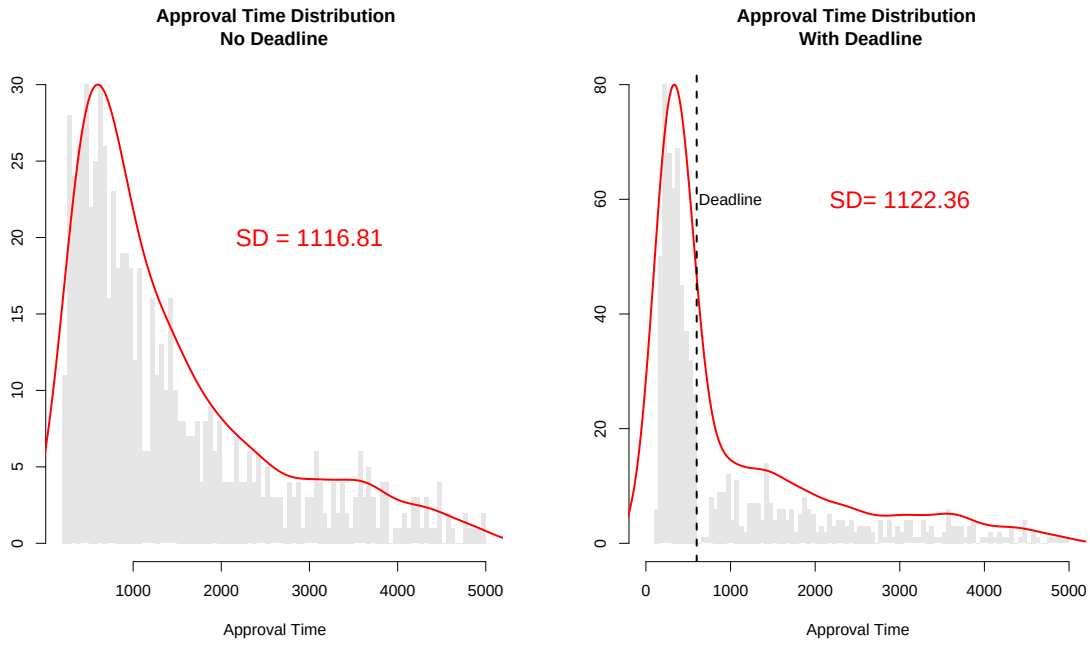


This figure demonstrates the conclusions of Lemma 2 and Proposition 4 deadlines induce a bimodal approval distribution and can *increase* the variance in the approval distribution. The left-hand plot shows the approval time distribution for a drug and the right-hand plot shows the same distribution, but now with a deadline imposed. Notice, that no drugs are approved immediately after the deadline inducing the bimodal distributions. Further, because the deadline bonus shifts a large proportion of the density before the bonus, we see that the variance actually increases due to the deadline bonus.

Figure 5: The Error Induced by Deadline Bonuses is Qualitatively Similar to the Error induced by Optimistic Priors on $\mu_i$

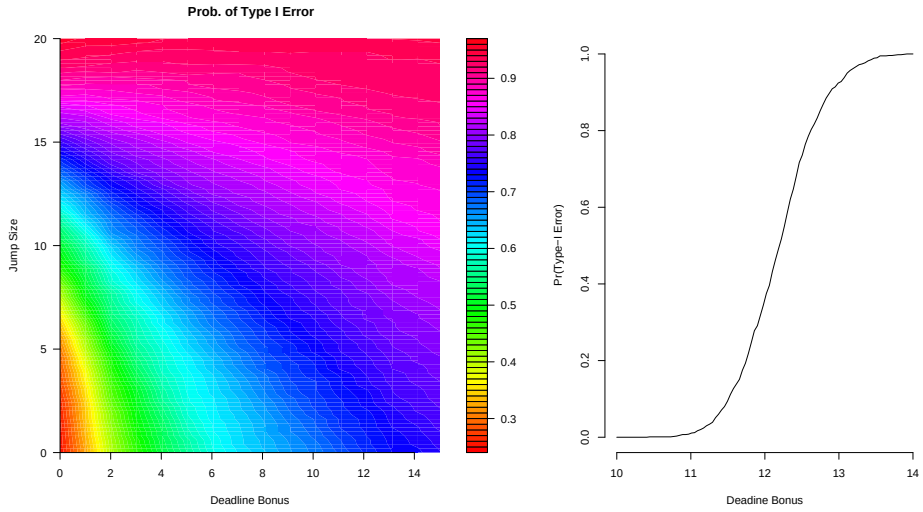


This figure presents the conclusions of Proposition 5. The left-hand figure shows that the probability of a type-I error (plotted along the vertical axis) is increasing as the deadline bonus increases (horizontal axis). The curve represents the type-I error rate over a set of simulations. The right-hand plot shows that the effects of deadlines on type-I error rates are mediated by the jump-size, with larger jumps causing deadlines to induce more error. This suggests that the optimal placement of a deadline is a highly non-linear problem dependent upon all the parameters considered here.

Figure 6: Deadlines Can Make Regulation Less Predictable

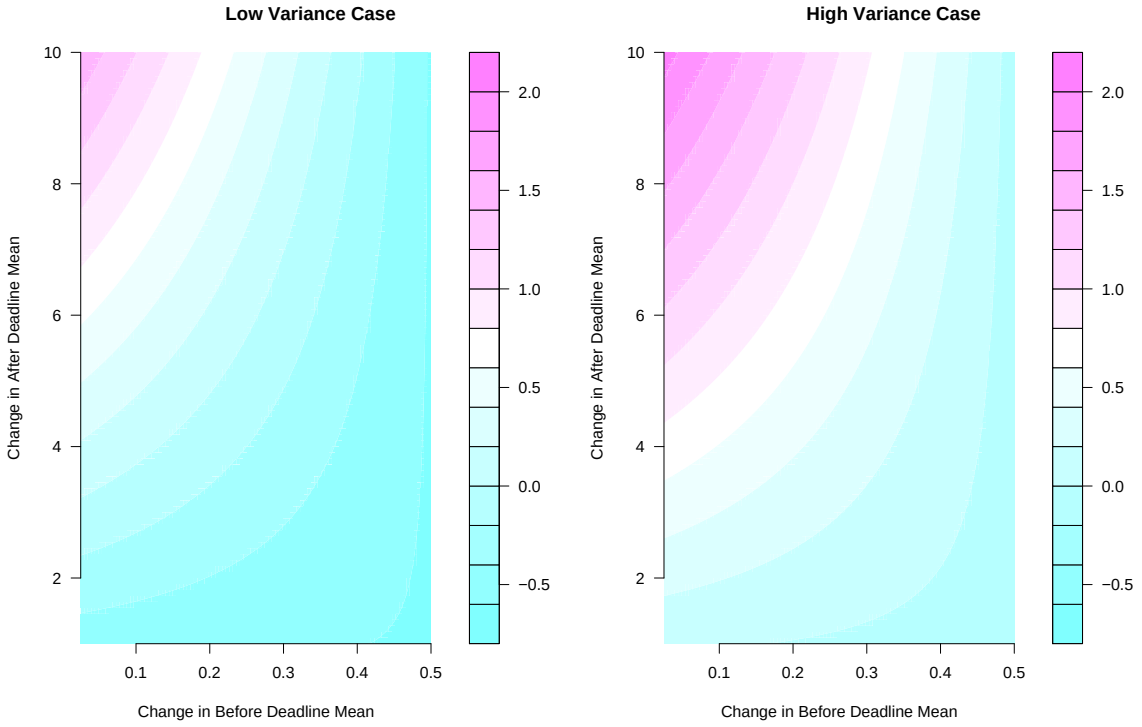


This figure evaluates Equation 6.2, applying different changes to the underlying mean and for a high- and low-variance case. This figure shows that deadlines can increase the variance in the approval distribution, depending exactly on how the deadline affects the regulatory behavior. Furthermore, this figure shows that the effect of deadlines on the variance in approval times is complicated, making it unlikely elected officials would be able to intentionally alter the predictability of drug reviews through the imposition of deadlines

Figure 7: Increases in Staff Do Not Translate into Decreased Variance in FDA Approval Times

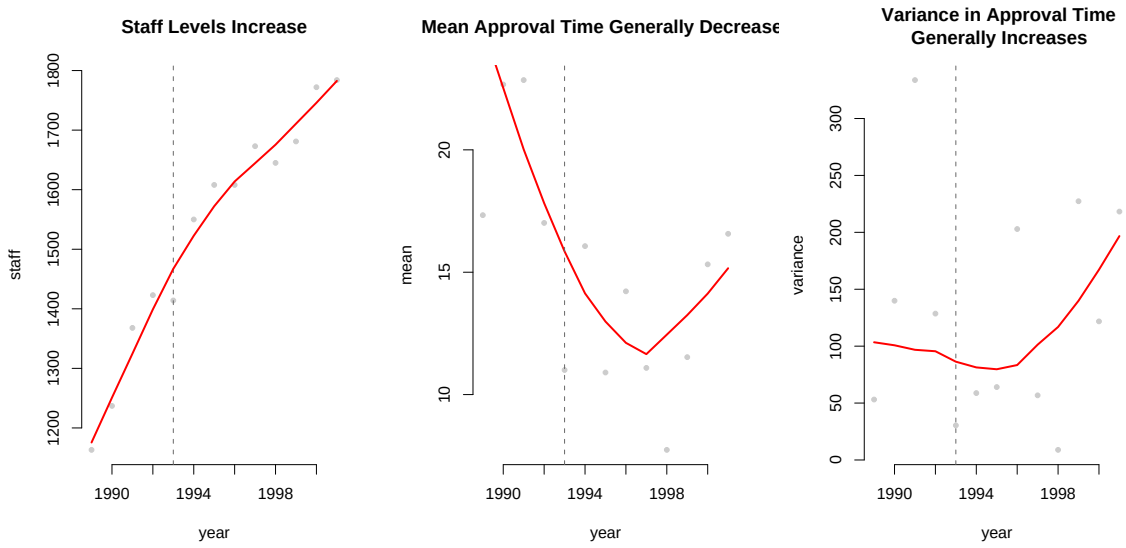


This figure demonstrates that the imposition of deadlines may affect the variance in approval times. The left-hand plot shows that during the 15 year period under consideration, the size of the staff for review exploded. This, coupled with the PDUFA reforms, is associated with the mean time-to approval in each year (vertical axis, center-plot) decreasing. But, contrary to intuition, the increases in staff do not affect the variance in approval times (right-hand plot, vertical axis). Rather, the variance in approval times is unchanged by the increases in staff. We suggest this could be attributed to deadlines.

Figure 8: Deadlines Affect Approval Time Distribution

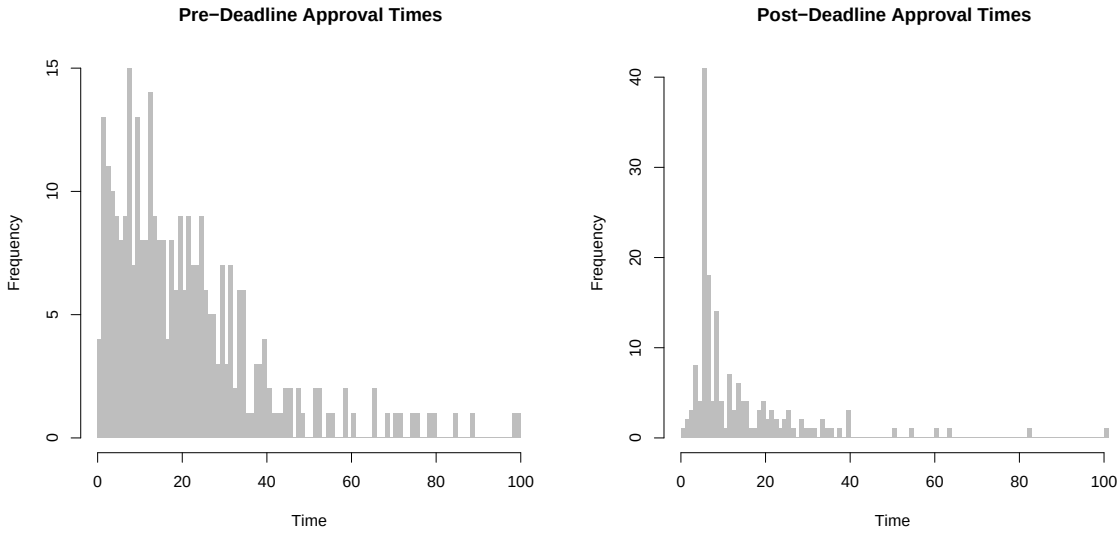


This figure shows the pre- and post-deadline approval time distributions. The left-hand plots shows that the approval time distribution before the implementation of PDUFA was unimodal and fairly smooth. By contrast, the right-hand plot shows that the PDUFA deadlines caused approval times to become multimodal, with approvals spiking immediately before the deadline elapsed. As argued above, this can cause variance to increase, potentially diluting the effects of increase staff on FDA approval times.