

Regulatory Fog: The Role of Information in Regulatory Persistence

Patrick L. Warren and Tom S. Wilkening*

September 26, 2012

Abstract

Regulation is very persistent, even when inefficient. We propose an explanation for regulatory persistence based on *regulatory fog*, the phenomenon by which regulation obscures information regarding the value of counterfactual policies. We construct a dynamic model of regulation in which the underlying need for regulation varies stochastically, and regulation undermines the social planner's ability to observe the state of the world. Compared to a full-information benchmark, regulation is highly persistent, often lasting indefinitely. Regulatory fog is robust to a broad range of partially informative policies and can be quite detrimental to social welfare. Regulatory experiments, modeled as costly and imperfect signals of the underlying state, do not eliminate the effects of regulatory fog. We characterize their effects and provide a framework for choosing amongst a set of potential regulatory experiments.

JEL Classification: L51, D82, D83

Keywords: Regulation, Information, Experimentation

*Patrick Warren: John E. Walker Department of Economics, Clemson University, 222 Sarrine Hall, Clemson, SC 29634. E-mail: pwarren@clemson.edu. Tom Wilkening: Department of Economics, The University of Melbourne, Economics and Commerce Building, Victoria, 3010, Australia. E-mail: Tom.Wilkening@unimelb.edu.au. We thank seminar participants at the Conference on Gaming Incentive Systems, USC's Marshall School, ISNIE 2010, The Australian National University, and the Clemson IO Workshop for useful comments. We want to specifically acknowledge Simon Loertscher, Kieron Meagher, Meg Meyer, Claudio Mezzetti, Dan Miller, Andy Newman, Harry J. Paarsch, Charles J. Thomas and Dan Wood for helpful suggestions and feedback.

1 Introduction

Many policies, particularly regulatory policies, alter what can be learned from observing the environment. A policy which successfully deters fraud, for instance, makes unscrupulous businessmen act like honest ones; price controls limit the ability of regulators to observe the consequences of more laissez-faire alternatives; and blanket ordinances limit the ability of individuals to express heterogeneous preferences. The objective of this paper is to explore the consequences of policies that, when put into place today, impact the information which can be used to select the policies tomorrow; what we refer to as *regulatory fog*.

Regulation in our model is a policy that pools or constrains different types into the same (second-best) action so that the state is unobservable. In contrast to many regulatory policies that can generate information, this type of regulation eliminates or distorts information about the state of the world. We believe that many regulations have this character, especially those such as simple price regulation, entry restrictions, and quotas that are very common and often designed without considering the informational consequences. Our main interest is in understanding the dynamics of optimal policy in cases where some policies reduce what can be learned about the underlying environment.

We study an infinite-horizon model with two Markovian states. In each period, a benevolent social planner has the option to regulate or not to regulate. Regulating is optimal in the bad state while not regulating is optimal in the good state. Regulatory fog arises because regulation put into place obscures the underlying state, whereas with no regulation, the state is observed. Therefore, at each point in time the planner faces a tradeoff between the benefit of the policy and the information this policy generates for future decision making.

We show that the slow and imperfect update process associated with regulatory fog can have significant consequences on the adaptability of policy and on welfare, relative to a full-information benchmark where both policies are informative. As regulation eliminates information about the underlying state, beliefs adjust monotonically toward the stationary distribution over time. As beliefs at this stationary distribution are often pessimistic, removing regulation has a high potential of leading to a deregulatory disaster. The potential of these disasters often swamp the positive informational advantages of removing regulation. For reasonable parameter values regulatory fog increases both the duration of an individual regulatory spell and the overall proportion of time spent under regulation. In cases where the planner is impatient, regulation may become permanent regardless of the persistence of

the underlying state.

Our baseline model reveals costs to regulatory policy that are not often taken into account in the debate over optimal policy but which are consistent with a number of empirical regularities that are difficult to rationalize with extant theory. First, due to an incentive to delay deregulation until the planner is reasonably sure that it will be successful, our model predicts large expected returns when detrimental regulation is finally lifted. This result is consistent with the estimated gains from deregulation of many network industries, including transportation, power, communications, and even banking.¹ Second, as deregulation is fraught with uncertainty, our model predicts that the removal of regulation can sometimes end in disaster, since it's a gamble every time.

One response to regulatory fog is to consider other sources of information. In many cases a planner can implement a range of smaller scale policy experiments or gain information exogenously from the actions or experiments of others. We model these “policy experiments” under the assumption that they are less risky than full deregulation but provide a weaker signal about the underlying state.² Using this framework we characterize the planner's induced preferences over regulatory experiments, including the optimal trade-off between costliness and effectiveness. We show that while policy experiments weakly reduce regulatory persistence, their value is non-linear in their effectiveness.³ This non-linearity reduces the value of imprecise signals and implies that experimentation is likely only when the experiment is informative and low cost. Our model suggests that many policy makers may optimally wait for external information to demonstrate the effects of deregulation rather than experiment on their own. Such “demonstration effects” are a feature of historical deregulation initiatives, as discussed below.

The difficulty in finding direct empirical evidence of regulation sustained by regulatory fog is self-evident. However, the role that external information shocks have played in historical deregulation suggests that a lack of information is a major deterrent to deregulation. The persistence of entry, price, and route regulation under the Civil Aeronautics Board (CAB) provides a useful example of this phenomenon. The CAB managed nearly every aspect of the airline industry, including fare levels, number of flights per route, entry into routes,

¹See, for example Moore (1988), Winston, Corsi, Grimm and Evans (1990), Peltzman and Winston (2000), and Joskow (2006).

²For example, an energy regulator may initially introduce competition for large commercial firms as a way to understand the potential benefits of broader deregulation. Such partial deregulation requires additional overhead and legal support but also generates partially informative information about the outcome of broader commercial and residential deregulation.

³Holding the cost of policy experiments fixed, a slight improvement in the informativeness of a policy which is initially uninformative yields little return to the policy maker. Vice versa, a slight improvement in the informativeness of a policy which is highly informative has a very large return to the policy maker.

entry into the industry and safety procedures from 1938 - 1978.⁴ Initially growing out the government's desire to regulate the pricing of air mail service, the CAB rules regulating passenger transport predate any substantial market for air travel. Despite the fact that the market for air travel essential grew to maturity under the umbrella of the CAB, the longevity of the regulatory body remains mysterious: the CAB generated an extremely inefficient set of economic regulations, as became apparent on their removal in 1978.⁵ How did such inefficient regulation persist, and why did it end when it did?

Critical to airline deregulation was the growth in intra-state flights, especially in Texas and California, because they revealed information about the likely effects of deregulation. These intra-state flights, and the local carriers who worked them, were not subject to regulation under CAB, so they gave consumers a window into what might happen if regulation was dropped more generally. A series of influential studies starting from Levine (1965) and continued and expanded by Jordan (1970) demonstrated that fares between San Francisco and Los Angeles were less than half the cost of those between Boston and Washington, D.C., despite the trips being comparable distances. Similar results were observed when looking at flights within Texas. There was also no discernable increase in riskiness, delay, or evidence of so-called "excessive competition."

The dissemination of these large-state market results proved to be a major catalyst for deregulation.⁶ The proximate driver of deregulation was a series of hearings held in 1975 by the Subcommittee on Administrative Practice and Procedure of the Senate Committee on the Judiciary (the so-called Kennedy Hearings). An entire day of testimony at these hearings was dedicated to exploring the comparison of intra-state and inter-state flights. William Jordan testified extensively, explaining and defending the results of the deregulatory studies.

The successful deregulation of airlines opened the door for deregulation in other related industries. The architect of the CAB deregulation, Alfred Kahn, cited the importance of the "demonstration effect," provided by airline deregulation, in understanding subsequent deregulation of trucking and railroads (Peltzman, Levine and Noll 1989). Likewise, the US experiment spurred airline deregulation overseas (Barrett 2008).

Consistent with our model of regulatory fog, a glimpse of the unregulated intra-state market provided regulators with new information about the underlying environment, which

⁴For an extensive review of the CAB's powers and practices, see "Oversight of Civil Aeronautics Board practices and procedures : hearings before the Subcommittee on Administrative Practice and Procedure of the Committee on the Judiciary", United States Senate, Ninety-fourth Congress, First session (1975).

⁵For an analysis of the economic effects of airline deregulation, see Morrison and Winston (1995) and Morrison (2007).

⁶Derthick and Quirk (1985) lay out the politics and timing of the push for deregulation, and cite these academic studies as the primary "ammunition" for those in favor of deregulation, as have others who have investigated the issue. See, for example, Bailey (1980) and Panzar (1980).

in turn spurred a successful deregulation of the airlines. The success of airline deregulation provided information which prompted deregulation in related industries.⁷

Finally, as our model would predict, deregulation often generates new information regarding the underlying primitives of an industry, as seen by the numerous academic studies which use the after-effects of deregulation to understand the composition of the industry. In addition to the long literature of airlines cited above, significant empirical analysis has occurred after the deregulation of trucking (Rose 1985), railroads (Boyer 1987), and cable television (Rubinovitz 1993), to cite just a few examples.

1.1 Related Literature

While we know of no paper which explicitly studies regulatory fog, our paper contributes to a broader literature interested in the causes of policy persistence. The dominant theory of policy persistence is political, in which rent seeking by entrenched groups is the primary force. Coate and Morris (1999) develop a model in which individuals make investments in order to benefit from a particular policy. Once these investments are made, the entrenched firms have an increased incentive to pressure the politician or regulator into maintaining the status quo. Similar dynamics can be found in Brainard and Verdier (1994), which studies political influence in an industry with declining infant industry protection.

As these models have no role for incomplete information, they have a hard time explaining the specific dynamics of regulation and deregulation. For us, one of the key features that distinguishes regulation from other policies is that it forces agents to take certain actions (or proscribes certain actions), and so generates similar signals in different states of nature. Contrast, for instance, the persistence of the CAB regulation to the huge variation in the US tax code over the same period (Piketty and Saez 2007). This effect is the essence of regulatory fog.

Asymmetric information has been combined with rent seeking models by Fernandez and Rodrik (1991). In their paper, uncertainty concerning the distribution of gains and losses of new legislation leads to lukewarm support by potential beneficiaries. Since uncertainty alters voting preferences in favor of the status quo, efficiency-enhancing legislation is often blocked by incumbents. In their model it is the aggregation of uncertainty across consumers which leads to persistence. By contrast, we find persistence naturally arising even in situations where a single planner maximizes social welfare. A recent paper by Friedrich (2012) also

⁷Given the Global Financial Crisis, some may argue that the deregulation wave went too far. Notice, however, that catastrophic deregulation failures are characteristic of regulatory fog. The inability to observe the counterfactual inevitably leads regulators to remove some regulation which is socially beneficial.

has a capture story, but with an interesting twist. In his model, negative net-present-value policies are continued, even by non-captured policymakers, in order to induce the advocates of positive net-present-value policies to put forth the costly effort to bring those projects to light. Captured policymakers can then use these good-intentioned policy extensions to hide their capture. Our model has no such incentive problem, yet still delivers persistence of bad policy from a non-captured social-planner.

A second extant explanation for policy persistence is that investment by firms leads to high or infinite transaction costs for changing policy. Pindyck (2000) calculates the optimal timing of environmental regulation in the presence of uncertain future outcomes and two sorts of irreversible action: sunk costs of environmental regulation and sunk benefits of avoided environmental degradation. Just as in our model, there are information benefits from being in an environment without regulation, and a social-welfare maximizing planner takes these benefits into account when designing a regulatory regime. Zhao and Kling (2003) extends this model to allow for costly changes in regulatory policy. Transaction costs act to slow changes in regulation, thereby creating a friction-based policy inertia. In our model, policy inertia is generated endogenously by the information that policies produce about the underlying state of the world. We attribute inaction by policy makers to their desire to wait for the environment to improve, which reduces the cost of experimentation and drives up the value of information.

While not directly related to policy persistence, our notion of regulatory fog is similar to parallel research being done in the context of social learning. Jehiel and Newman (2011) study an intergenerational environment in which contracts put into place today limit the observation of potentially detrimental actions in the future. As principals are replaced each generation, the suppression of information from past contracts leads to the existence of loopholes where contracts offered by future principals are exploited by unconstrained agents. While our model reveals that inefficient policies can be generated in a dynamic setting with perfect recall, their model shows that similar properties can hold in a static environment when policy makers have limited recall.⁸ Peck and Yang (2010) study information flows in an environment with social learning and investment flows which change according to an underlying Markov process. Their model reveals that the ability of agents to delay action until they observe the actions of others can lead to asymmetric investment cycles. While we do not explicitly discuss the case of multiple jurisdictions, the persistence of policies which eliminate information are similarly exacerbated in our framework when multiple jurisdictions and social learning are introduced. Finally, our model shares some features with Ichino and

⁸For other models of policy persistence with nearly-myopic policy makers see Berentsen, Bruegger and Loertscher (2008) and McLennan (1984).

Muehlheusser (2008), which examines the optimal oversight of a subordinate whose type is unknown. As in our model, there is a trade-off between information acquisition and short-run incentives.

In the next section, we lay out the basic model, and solve for the optimal regulatory policy under both the full-information benchmark and a simple incomplete information environment. In section 3, we compare the results of these two models to illustrate the effects of regulatory fog on persistence and welfare. In section 4, we expand the policy space and evaluate small-scale deregulatory experiments with intermediate information consequences. Section 5 concludes.

2 Model

Consider an economy which is home to a large number of consumers and producers, both potential and actual, and a welfare-maximizing social planner who lives forever and is risk and loss neutral.

There are two possible states of the world, good (G) and bad (B). In the good state, the free-market equilibrium is efficient, delivering some level of social welfare, which we normalize to zero. In the bad state, some imperfection inhibits efficiency in the free-market equilibrium. There are many ways to think about this imperfection: technology has shifted in a way that gives rise to a natural monopolist; a new mode of production has been invented which is profitable but involves a pollution externality; or network characteristics of the product are improved with standardization. Regardless of the origin of the imperfection, society suffers a deadweight loss when the free market reigns in the bad state that we normalize to -1 . No individual member of society knows the underlying state, including the planner. Rather, each responds to his localized incentives, which may be quite idiosyncratic, and the state only becomes apparent when the market outcome is observed.

The planner can enforce a regulation that eliminates the deadweight loss arising from the imperfection, ($R \in \{0, 1\}$), where regulation entails its own economic cost $-d$. This regulation might be the establishment of a price-capped regulated monopoly, a cap on pollution output enforced with auditing and fines, or a blanket ordinance eliminating discretion. We assume that $1 > d$ and thus the planner will prefer to pay the cost of regulation rather than simply accept the loss from having an unregulated market in the bad state. We also assume that the cost of regulation does not depend on the underlying state, so the total welfare under regulation is $-d$, regardless of the state. The planner has a discount factor of δ , and

will regulate if indifferent.⁹

Transitions between the two states follow a Markov process where ρ_{ij} represents the probability of transitioning to state j from state i between two periods. Both the good and the bad states are persistent with $\rho_{BB} \in (.5, 1)$ and $\rho_{GG} \in (.5, 1)$. The transition probabilities are known to the planner.

The timing of the model is as follows: at the beginning of every period, the planner chooses the policy environment R . Next, nature chooses the state according to the probabilities above.¹⁰ Finally, the market responds to the policy choice and the state.

The per-period value to the planner for regulating and not regulating in each state is given by:

$$(1) \quad \begin{array}{rcc} & \begin{array}{cc} \textit{Good State} & \textit{Bad State} \end{array} \\ \begin{array}{l} \textit{Regulation} \\ \textit{No Regulation} \end{array} & \begin{array}{cc} -d & -d \\ 0 & -1 \end{array} \end{array} .$$

While the planner would prefer to regulate in the bad state and to not regulate in the good state, the current period's regulation decision alters the information available to the planner for future decisions. Under regulation the planner gains no new information about the underlying state, as both states look identical, and simply updates according to the transition probabilities and her prior belief. In the absence of regulation, she learns the state for certain, since if the state is bad, imperfections will arise and lower social welfare.¹¹ This difference in information generated by different policy choices is the information cost we explore throughout the paper.

In the base model, we make the information effect of regulation extremely stark, since the key intuitions are apparent in that setting. In reality, some information about the

⁹We assume here that the planner is strictly publicly-interested. Allowing for some degree of interest-group oriented regulation in the spirit of Stigler (1971) or Grossman and Helpman (1994), does not substantively change the underlying persistence we are exploring. While the particulars of the planner's objective function will affect the relative value of different states and the interpretation of actions, the impact of regulatory fog on regulation and efficiency are quite similar.

¹⁰Note the difference between this model and the K -arm bandits originally developed by Robbins (1952). A distinguishing feature of bandit problems is that the distribution of returns from one arm only changes when that arm is chosen. This feature implies that the distribution of returns does not depend explicitly on calendar time and there is no new information being generated about other actions. By contrast, a fundamental decision in our model is the potential improvement of alternative actions through waiting. See Bergemann and Välimäki (2006) for a discussion of bandits.

¹¹We realize that deregulation is, in reality, a slow process, the full consequences of which take can take many years to ascertain (Winston 1998). We make the simplest assumption of immediate revelation in order to make the contrast of interest, between regulatory fog and the full-information benchmark, as transparent as possible. Since the assumption about the consequences of deregulation will be held fixed across all these comparisons, we chose simplicity over realism in this dimension of our model.

state is available even under regulation, but less extreme informational differences generate substantively similar results. We consider the case where some information about the state is revealed under regulation in section 4.

In the next two subsections, we solve for the planner’s optimal policy choice in two environments. First, we will consider the optimal way to set a benchmark policy that, unlike regulation, does not affect the observability of the state of the world. Next, we will consider regulation as described above. Our goals are to contrast the planner’s choices with these two types of policy, derive the implied difference in policy persistence, and calculate the welfare consequences of regulatory fog. We are not trying to compare regulation to deregulation, but rather to contrast the two activist policies, one with regulatory fog and one without.

2.1 Full Information Benchmark

Before developing the optimal use of regulation, it is useful to consider the planner’s optimal policy when information is unaffected by market intervention. For ease, we refer to policies which do not affect information as non-regulatory policy *interventions*.

Consider a small change to the model above, in which the planner observes the state at the end of every period, regardless of the policy decision. We determine the optimal use of non-regulatory interventions, and use this as a benchmark for identifying the effects of regulatory fog.

As the planner knows the previous period’s state with certainty, the information environment is greatly simplified. If the state was good last period, the probability that the state is bad this period is given by $\rho_{GB} < 0.5$. Likewise, if the state was bad the probability that the state remains bad is $\rho_{BB} > 0.5$. The planner is not clairvoyant, as she does not observe the state before she makes her policy decision for the period, but she does learn what the state was at the end of the period, even if she chooses to intervene. The following proposition characterizes the policy function of an optimal planner:

Proposition 1 *Assume that the state is revealed at the end of each period. Then the planner’s optimal policy falls into one of the following cases:*

1. *If $d \leq \rho_{GB}$, the planner intervenes every period.*
2. *If $\rho_{GB} < d \leq \rho_{BB}$, the planner intervenes after the bad state and does not intervene after the good state. Conditional on enactment, the length of a intervention follows a geometric distribution with expected length $1/\rho_{BG}$. The proportion of time spent with*

government intervention is given by the steady state probability of the Markov Process $\left(\frac{\rho_{GB}}{\rho_{GB}+\rho_{BG}}\right)$.

3. If $d > \rho_{BB}$, the planner never intervenes.

Proof. All proofs in the appendix. ■

Proposition 1 identifies the key comparisons that drive the planner’s decision when regulatory fog is not a problem. Recall that d is the cost of intervening which scales between zero and the cost incurred to society in the unregulated bad state, which we normalized to -1. When d is small relative to the probability of transition to the bad state, the planner will intervene in every period regardless of last period’s state. This permanent intervention reflects the rather innocuous costs of intervention relative to the potential catastrophe of being wrong.

Likewise, if the cost of intervention is very high relative to the cost of the unregulated-market imperfections, the planner prefers to take her chances and hope that the underlying state improves in the next period. The interventionist cure is, in expectation, worse than the disease and leads to a *laissez-faire* policy.

The interesting case for our model is the range of intermediate costs for which the planner finds it in her interest to adapt her policy to the information generated in the previous period. In the full information case, the planner applies the policy which is optimal relative to the state observed in the previous period. Except for the periods in which the state actually transitions, the policy adopted by the planner will be *ex post* second-best efficient. Even in this full-information environment, there is some policy persistence. If the planner intervenes this period she is more likely to intervene next period, since the underlying state is persistent.

2.2 Optimal Regulation with Regulatory Fog

We now return to the base model and solve for the planner’s optimal policy when the regulated state eliminates information. We construct this policy in three steps: we first define the process by which the belief of being in the bad state is updated and construct the value functions of regulating and not regulating as a function of that belief. Next, noting the regularity in these value functions, we show that there is a unique cutoff belief for which the planner switches from regulation to deregulation. We then use this insight to construct the optimal policy function for the planner under regulatory fog.

Just as in the full-information benchmark, the planner’s decision to regulate and deregulate in each period is essential a tradeoff between (i) regulating today and suffering a certain

payoff loss of $-d$ or (ii) choosing not to regulate today and receiving either 0 or -1 depending on the underlying state. As the value of the lottery is based on the likelihood of being in the bad state, a planner's decision to regulate will be guided by his belief regarding the state of nature in each point in time.

Unlike the full-information benchmark, however, the evolution of this belief is guided by the policy chosen in each period. A planner who does not regulate immediately learns the state of nature and thus has beliefs which are very optimistic or pessimistic regarding next period's underlying state. A planner who regulates, by contrast, learns very little about the underlying state of nature and updates her beliefs only in relation to the underlying transition matrix. The removal of regulation generates information which is valuable to the planner for future decisions. The tradeoff between the expected gains and losses of the current period and the information gained by deregulation form the basis of decision making for the planner over time.

Let ϵ be the belief of the planner that he is in the bad state today. Further, define an updating function $P \rightarrow [0, 1]$ which maps the current belief into a posterior belief in the case where regulation is imposed. Using the underlying transition matrix,

$$(2) \quad P(\epsilon) = \epsilon\rho_{BB} + (1 - \epsilon)\rho_{GB}.$$

Let $P^k()$ represent k applications of this function, where negative numbers represent inversions. Then for any starting $\epsilon \in [0, 1]$,

$$\lim_{k \rightarrow \infty} P^k(\epsilon) \equiv \tilde{\epsilon} = \frac{\rho_{GB}}{\rho_{GB} + \rho_{BG}}.$$

Further note that $P(\epsilon)$ is continuous and increasing in ϵ , $P(\epsilon) \leq \epsilon$ for $\epsilon \geq \tilde{\epsilon}$, and $P(\epsilon) \geq \epsilon$ for $\epsilon \leq \tilde{\epsilon}$. These conditions imply that the uninformed planner updates over time toward the steady state of the Markov process. This implies that the potential posteriors of no regulation bracket the posterior of regulating.¹²

Let $R(\epsilon) \in \{0, 1\}$ represent the planner's decision when she believes the state is bad with probability ϵ , where $R = 1$ indicates regulation and $R = 0$ indicates no regulation. Let $V(R|\epsilon)$ be the planner's value function playing regulatory policy R with beliefs ϵ . Define $R^*(\epsilon)$ as the maximizing policy function and $V^*(\epsilon)$ as the value function induced by it.

¹²In the language of informational decision analysis, one posterior brackets another if for any signal, all posteriors of one information system can be written as a linear combination of the posteriors of the other information system with linear weights less than one. Bracketing ensures that the expected value of information from removing regulation is weakly positive. See Hirshleifer and Riley (1992).

Given maximization in all subsequent periods for any belief ϵ ,

$$(3) \quad V(R = 1|\epsilon) = -d + \delta V^*(P(\epsilon)),$$

$$(4) \quad V(R = 0|\epsilon) = \epsilon[-1 + \delta V^*(P(1))] + (1 - \epsilon)[\delta V^*(P(0))].$$

For notational simplicity, let $V_B \equiv V^*(P(1))$ and $V_G \equiv V^*(P(0))$. V_B represents the value function after observing a bad state while V_G represents the value function after observing a good state.

A useful way of interpreting these value functions is by taking their difference. Define

$$(5) \quad G(\epsilon) = V(R = 1|\epsilon) - V(R = 0|\epsilon)$$

as the difference in value between regulation and deregulation given beliefs ϵ . Substituting from (3) and (4) in equation (5) yields:

$$G(\epsilon) = \underbrace{\epsilon - d}_{\substack{\text{Expected} \\ \text{Cost of} \\ \text{Deregulation}}} - \underbrace{\delta[\epsilon V_B + (1 - \epsilon)V_G - V^*(P(\epsilon))]}_{\text{Expected Value of Information}}.$$

The first term represents the expected current period cost of deregulating, since the planner will suffer the bad state with probability ϵ , but saves the cost of regulation (d). The second term represents the value of information associated with learning the true state: instead of having to work with a best guess of $P(\epsilon)$, the planner will know with certainty that she is in the good or bad state and can act accordingly.

Figure 1 shows the current period cost of deregulation and the value of information over the domain of ϵ . The expected cost of deregulation is linear; negative when $\epsilon = 0$, and positive at $\epsilon = 1$. By contrast, the value of information is concave and equal to zero at both endpoints. It follows directly that there exists a unique point where $G(\epsilon) = 0$.

Proposition 2 *There exists a unique cutoff belief $\epsilon^* \in [0, 1]$ such that the optimal policy for the planner is to regulate when $\epsilon > \epsilon^*$ and to not regulate when $\epsilon < \epsilon^*$.*

In this dynamic setting, the value of information relates strongly to the static models of Hirshleifer and Riley (1979) and Radner and Stiglitz (1984): information is most valuable when the planner is least certain about the underlying state. At $\epsilon = 0$ and $\epsilon = 1$ the planner knows the underlying state and thus learns no new information by deregulating. In these cases, the value of information is zero. In the interior, the value of information is strictly

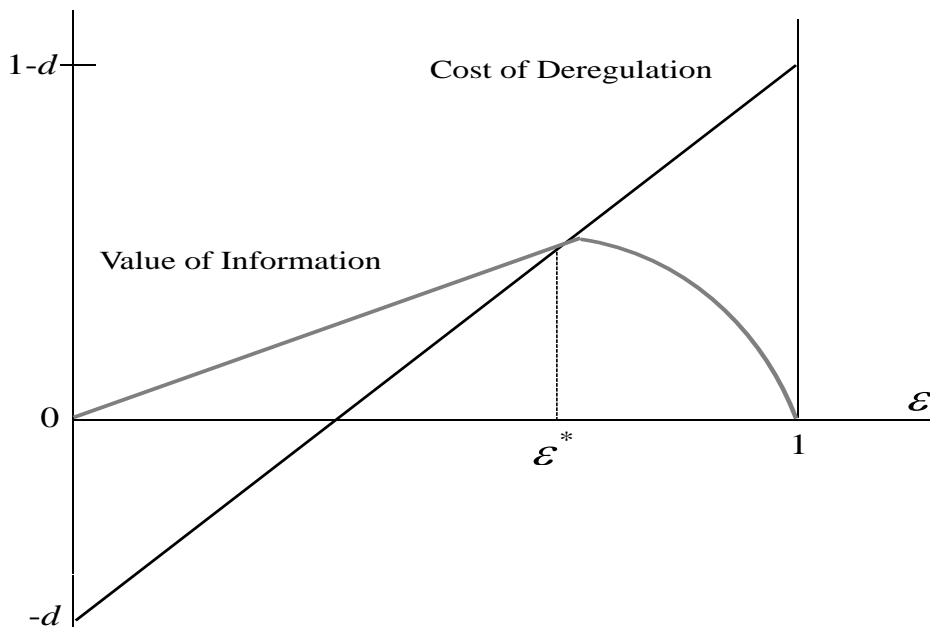


Figure 1: The cost and value of deregulation: Uniqueness of ϵ^*

positive and concave.¹³

Proposition 2 provides structure to the solution of the planner's problem. Although the optimal regulation decision is defined for any belief ϵ , only countably many (and often finite) beliefs will arrive in equilibrium. Let ϵ^* be the planner's optimal cutoff as defined in Proposition 2, and define k^* as the unique $k \in \mathbb{N}^*$ such that $P^{k+1}(1) \leq \epsilon^* \leq P^k(1)$. We refer to k^* as the length of the regulatory spell. If there does not exist a k which satisfies this condition, then $k^* = \infty$. This will be the case if and only if $\epsilon^* \leq \tilde{\epsilon}$.

As with the full information benchmark, our goal is to relate the proportion of time spent under regulation to the cost of regulation d . As the length of regulatory spells (k^*) is a weakly decreasing function of ϵ^* , it is useful to first determine how ϵ^* changes with respect to d .

Corollary 1 *The threshold ϵ^* is increasing in d .*

The intuition for Corollary 1 can be seen in Figure 1. As d increases, the direct cost of deregulation decreases. As a result, the cost curve shifts downward, which shifts ϵ^* to the

¹³The cone-like shape is due to differences in the update operator for beliefs above and below ϵ^* . For $P(\epsilon) < \epsilon^*$, a planner who does not deregulate this period will, optimally, deregulate in the next period. As the value function for no regulation is linear, the value of information increases linearly in this region. For $P(\epsilon) > \epsilon^*$ the planner has an incentive to maintain regulation for at least one period. The longer the delay before deregulation, the lower the expected cost of deregulating. Thus, the value of information decreases non-linearly in this domain due to the recursive nature of the updating operator $P^k(\cdot)$.

right. At the same time, an increase in d increases the cost of regulating, increasing the difference in payoffs between regulating and deregulating in the good state. This additional cost of regulation increases the value of information for all $\epsilon \in (0, 1)$ leading the value of information curve to expand upward. As both of these effects makes $G(\epsilon)$ smaller, the overall effect is an unambiguous increase in the deregulation cutoff.

As k^* is a weakly decreasing function of ϵ^* it follows:

Corollary 2 *The length of a regulatory spell, k^* , is weakly decreasing in d .*

Bringing these effects together, we can characterize the planner's optimal policy choice.

Proposition 3 *There exists a unique optimal policy function for the planner under regulatory fog. Let*

$$(6) \quad \tau \equiv \left[\frac{\delta + \frac{1-\delta}{\rho_{GB} + \rho_{BG}}}{1 - \delta[\tilde{\epsilon} - \rho_{GB}]} \right] > 1.$$

Once regulation is applied the first time the planner's optimal policy falls into one of the following cases:

1. *If $d \leq \rho_{GB}\tau$: the planner always regulates.*
2. *If $d > \rho_{BB}$: the planner never regulates.*
3. *If $\rho_{GB}\tau < d \leq \rho_{BB}$: Let ϵ^* be the solution to the implicit function:*

$$(7) \quad \epsilon - d + \delta[\epsilon V_B + (1 - \epsilon)V_G - V^*(P(\epsilon))] = 0.$$

The planner regulates for $k^ > 0$ periods after the bad state is revealed and does not regulate after the good state is revealed, where k^* is the first k such that $P^{k+1}(1) \leq \epsilon^* \leq P^k(1)$.*

The proof for this proposition is provided in the appendix, but the intuition is straightforward. Consider a planner that just experienced the bad state. She is at her most pessimistic at the beginning of the subsequent period, believing the state is still bad with probability ρ_{BB} . While regulating, the planner's belief about the underlying state trend downward, since $P(\epsilon) < \epsilon$. If the belief crosses the threshold defined in Proposition 2, she deregulates. If having no regulation reveals a good state, she remains without regulation until she experiences the bad state, at which point the process begins again.

For very costly regulation, ($d > \rho_{BB}$), even the most pessimistic beliefs the planner will ever have are not sufficient to justify regulation. For very low cost regulation ($d < \tau\rho_{GB}$), the most optimistic belief the planner ever has ($\tilde{\epsilon}$, the steady state of the Markov process) is not sufficiently rosy to justify the risk of deregulation. For intermediate costs, regulation will progress in cycles.

Having characterized the planner's strategy under regulatory fog, the next section compares the optimal regulatory policy to that in the full-information benchmark and performs a number of comparative-static exercises.

3 Comparison with Benchmark

Regulatory fog has two fundamental consequences in our model, and each affects both the time under regulation and overall social welfare. First, once regulation is imposed, beliefs under regulation always remain above $\tilde{\epsilon}$. This contrasts markedly with the full-information planner, who updates to the more optimistic ρ_{GB} after observing a good state, even while regulating. A decision maker considering whether to deregulate is always faced with the potential of a deregulatory disaster, wherein the removal of regulation in the bad state leads to losses. The chance of this disaster is higher with regulatory fog, no matter how long the planner waits, because it is bounded below by the steady state of the Markov process. This potential for disaster can lead to permanent persistence, particularly in environments where the decision maker is relatively impatient.

Second, the planner's belief about the underlying state under regulation evolves smoothly over time from a belief in which regulation is (almost) certainly optimal to a belief in which there is a greater likelihood that regulation is inefficient. This contrasts to the full-information planner whose beliefs bounce quickly from very optimistic to very pessimistic as a function of the observed state. For most cases this process naturally leads to regulatory inertia since delay (i) reduces the chance of deregulatory disasters and (ii) increases the value of information from deregulating.

We will discuss these effects in turn.

3.1 Permanently Persistent Regulation

We begin by studying the range of parameters for which regulation persists indefinitely. As with the full-information benchmark, regulation is fully persistent if the normalized cost is low relative to the probability of transition from the good to the bad state. However, as the

most optimistic beliefs that arrive in equilibrium are more pessimistic, deregulation carries additional risk, which is represented by τ in Proposition 3. Since $\tau > 1$, regulatory fog strictly increases the set of parameters for which regulation persists permanently.

As τ is a decreasing function of the discount factor δ , impatient planners are more affected by regulatory fog. Planners who discount the future completely ignore the value of information from deregulation and are willing to deregulate only if the probability of being in the bad state falls below the cost of regulation. Institutions that induce short-sighted preferences by regulators, such as having short terms in office, are expected to lead to more regulatory persistence. Consistent with this prediction, Smith (1982) finds that states with legislators having longer terms are more likely to deregulate the licensure of professions while Leaver (2009) shows that electricity regulators with longer terms will review and adjust rates more frequently than their short-term counterparts.¹⁴ Our example involving CAB deregulation is also consistent with this result: the initial push for deregulation came from Senator Kennedy, a long-termed senator with a very safe seat. Kennedy's secure tenure and expected longevity implies that he would enjoy the fruits of a successful deregulation but had limited personal costs if such deregulation failed.

Continuing with comparative statics, as the persistence of states increases, the effect of regulatory fog becomes more pronounced. Consider a proportional decrease in ρ_{BG} and ρ_{GB} . In the absence of regulatory fog, such a change decreases the range of costs for which permanently persistent policy occurs, since the observation of a good state this period becomes a much better signal about that state being good tomorrow. In the case of regulatory fog, such a change leaves the steady state $\tilde{\epsilon}$ unaffected, but increases τ . The net effect on the cost cutoff for permanently persistent regulation is ambiguous. However since τ grows as ρ_{GB} falls, the gap between the full-information benchmark and regulatory fog unambiguously increases. Empirically, this means we should expect to find examples of permanent persistence induced by regulatory fog in instances where the underlying need for regulation changes quite slowly, relative to the speed at which regulation can be changed, *ceteris parabis*. This might mean, for example, that regulatory fog is less of an issue for a very new industry in which the ebbs and flows of market power are still quite rapid, but more of an issue in an established industry (such as airlines) in which the technological conditions that induce market power change at a much slower pace.

Permanent regulation exists under regulatory fog even as the underlying states become highly persistent. Figure 2 shows the region of permanent regulation both for the case of discretely positive transition probabilities and for the case where ρ_{BG} and ρ_{GB} converge to

¹⁴For another example of technical and information constraints inducing myopic policy, see Aidt and Dutta (2007). Both Leaver (2009) and Aidt and Dutta (2007) consider a very different mechanism than ours.

zero, but where $\tilde{\epsilon} \in (0, 1)$. In the full-information case, planners always have an incentive to deregulate in the good state as ρ_{BG} and ρ_{GB} converge to zero. In the presence of regulatory fog, however, the most optimistic belief achievable in equilibrium is $\tilde{\epsilon}$, which may be quite pessimistic in the limit. Referring back to Proposition 3, $\tau\rho_{GB} \rightarrow 0$ as $\rho_{GB} \rightarrow 0$ if and only if $\delta = 1$. Otherwise it is bounded away from zero, and so for low costs, regulation will persist indefinitely, even though one deregulatory episode could lead to the (near) permanent removal of regulation.

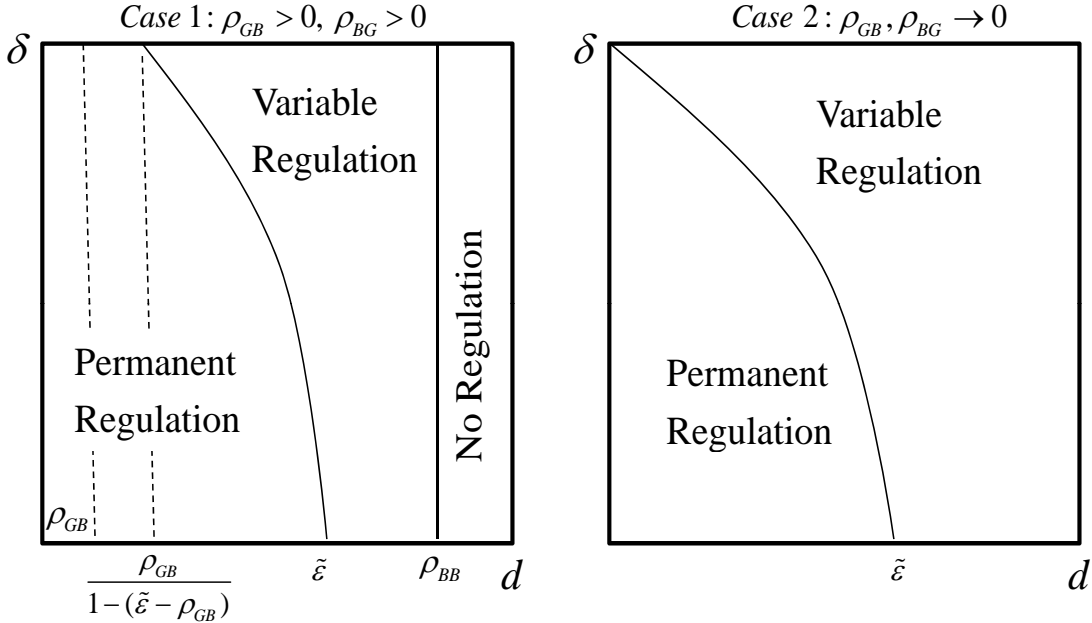


Figure 2: Optimal policy under regulatory fog as a function of impatience (δ) and the cost of regulation (d).

Finally, as illustrated by the CAB example in the introduction, there is a way to escape from the world permanently persistent regulation. If some information arrives from outside the system that leads the policymaker to update her beliefs, she could end up with beliefs that are more optimistic than the ϵ^* , even when that cutoff is below steady-state belief. We analyze the effects of the exogenous arrival of information in Section 4, below, as a special case of experimentation.

3.2 Regulatory Cycles

When the costs of regulation d are moderate, regulatory policy is characterized by transitions between regulation and deregulation, and these transitions are influenced by the underlying state.

As noted in Proposition 1, the transition from regulation to deregulation in the full information benchmark is based on the arrival time of the first good event and thus there is a direct relationship between persistence and the stochastic nature of the environment. As arrival times follow a geometric distribution, the expected length of a regulatory spell is $\frac{1}{\rho_{BG}}$, and the expected time under regulation is equal to the steady state probability $\tilde{\epsilon}$. Furthermore, for $d \in (\rho_{GB}, \rho_{BB})$, there is no relation between the cost of regulation and its persistence.

Unlike the full information case, regulatory fog is characterized by (often long) fixed periods of regulation followed by deregulation. Deregulation lasts until the arrival of the first bad event, at which point the regulatory cycle repeats. The overall effect of these cycles on policy can best be seen by plotting the proportion of time spent in regulation and deregulation as a function of d . As can be seen in Figure 3 regulatory fog leads to more persistence for small and medium d , and less persistence for large d . This differential effect is driven by the planner's trade-off under regulatory fog between (i) the potential negative outcome from deregulating in the bad state and (ii) the information learned about the underlying state, which can benefit future decisions. Under the full-information benchmark, there is no such trade-off.

This leads to quite different testable implications about the relationship between the cost of regulation and the length of regulatory spells in the presence and absence of regulatory fog. In the presence of regulatory fog, the length of regulatory spells should shrink as the cost of regulation increases, but the likelihood of failed deregulation should rise. In the absence of fog, there is no relation between the cost of regulation and the length of regulatory spells on the likelihood of failed deregulation (except at the extremes of permanent regulation and no regulation.)

When d is small, the relative cost of deregulating in the bad state is large, leading to delayed deregulation in order to reduce the chance for a deregulatory disaster. As d grows, the value of being in the deregulatory good state grows, while the additional cost to deregulation shrinks.¹⁵ The decline in persistence does *not* mean that regulatory fog is less important in these circumstances. In fact under regulatory fog with high regulatory costs, the planner simply replaces some of the time spent under regulation in the full-information environment with time spent in the unregulated bad state. As many deregulatory episodes are immediate

¹⁵In economic environments, we view the region of parameters for which rapid cycles of deregulation and regulation should occur to be quite rare. It is our view that planner myopia and moderate to low regulation costs are typically the norm. In other fields such as medicine, however, there is suggestive evidence that both regions exist. Treatment for cancer, for instance, is characterized by cycles in treatment and careful monitoring. On the other hand, treatment for high blood pressure or depression are continuous with little variation in treatment over time.

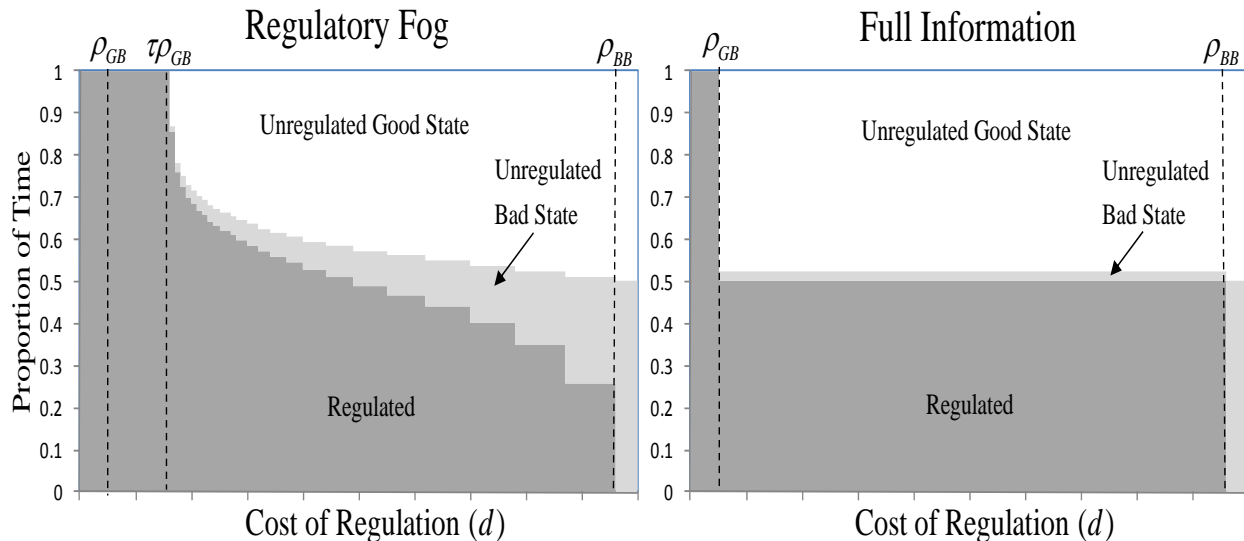


Figure 3: Proportion of Time in Each Regulatory State as a Function of The Cost of Regulation (d) and the Information Environment. Graph shown for $\rho_{GB} = \rho_{BG} = .05, \delta = .9$.

failures, overall social welfare decreases.

The planner's equilibrium payoffs in the full-information benchmark and under regulatory fog are presented in Figure 4. When the cost of regulation is very high or low, information has no value, since either regulation will always or never be applied. In these cases there is no cost of regulatory fog. Otherwise it imposes an information cost on the planner which is linear up to $\tau\rho_{GB}$ and concave thereafter. Overall, welfare loss is greatest for intermediate values of d where there is both large amounts of policy persistence and high amounts of failed experimentation.

We do not believe the CAB is a good example of the cycling phenomenon, but there is some evidence of historical cycles in electricity generation deregulation. While regulation of electricity generation in the 20th century was done primarily through the granting of state-sanctioned monopolies, almost half of the U.S. states had experimented with some form of competition by the year 2000.¹⁶ By 2006, however, only 12 states still had electricity deregulation. The rest either repealed their deregulatory policies or delayed their implementations. Norway, Sweden, and Canada had similarly tumultuous experiences, and have reintroduced regulation to some degree.¹⁷

¹⁶http://www.ftc.gov/bcp/workshops/energymarkets/background/slocum_dereg.pdf

¹⁷See Joskow (2008).

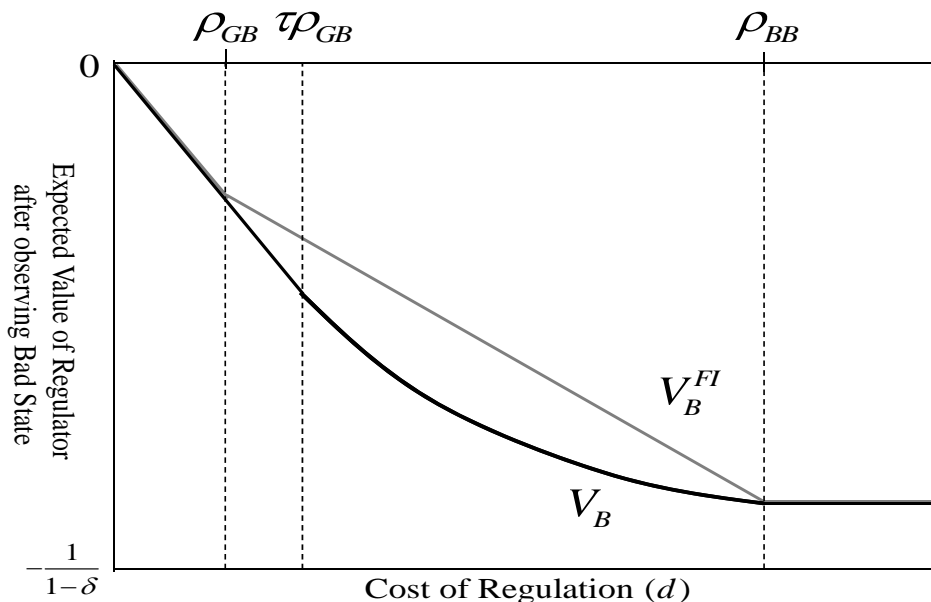


Figure 4: The Cost of Regulatory Fog ($\rho_{GB} = \rho_{BG} = .05, \delta = .9$)

4 Policy Experiments

Just how bad a problem is regulatory fog? In the preceding sections we have left the planner with the stark choice between full regulation and full deregulation and shown that, in such a world, regulatory fog leads to persistent regulation and significant welfare losses, relative to a benchmark policy that has no information consequences. We might wonder, however, just how bad the information problem is in an environment with a broader policy space or if some information may leak through even in a relatively strict regulatory regime. Furthermore, we may wonder why a planner cannot make small alterations to regulatory policy to generate new information without suffering the potentially disastrous consequences of full deregulation in the bad state. This section studies the planner's optimal policy when she has access to experimentation, a broader set of policy options that may be less efficient than full regulation but which are potentially more informative.

The experiments we consider in this section vary from deregulation in two ways. First, experimentation can be conducted while maintaining regulation, but these experiments may have an additional cost which is borne by society. These costs reflect both the direct overhead costs of measurement and the indirect costs of implementing mechanisms which may be more complex than the simple static regulation.

Informative mechanisms will often be very different from the static mechanism, and thus

the indirect costs of experimentation are unlikely to be trivial. In the case of regulation which reduces moral hazard, for instance, small changes in the incentive scheme can lead to a large change in actions, and is typically tantamount to deregulation. Thus in this case, a regulatory experiment which maintains regulation broadly must be more complicated than simply cutting back on the degree of monitoring. In a broader context, dynamic mechanisms will typically involve screening mechanisms, which must distribute information rents to a subset of the population or encourage other forms of inefficiency.

The second difference from full deregulation is the *imprecise* information attained from the small scale experiment about the underlying state. This imprecision comes from two sources. First, there are basic statistical problems associated with sampling a small selection of firms or markets. Even a perfect and unbiased experiment will have some sampling variance. There is also a risk that an improperly designed experiment may lead to spurious results. Second, the very circumscribed nature of the experiment may limit its usefulness. If firms expect the experiment to be temporary, for example, they may react very differently from how they would respond to a deregulation of indefinite length. The partial equilibrium response of agents to a deregulatory experiment may be very different from the general equilibrium response which would result from full deregulation.

To illustrate this idea, consider a planner who wants to know the probable effects of a general lowering of immigration restrictions and experiments by relaxing the immigration restriction to certain regions. Her experiment may give biased results for many reasons. If the demand for entry to the areas chosen was not representative of overall demand, she may under- or over-estimate the demand for entry. More importantly, the demand for entry to the selected regions may be directly affected by the partial nature of the experiment. If it is known to be a temporary loosening, immigrants may quicken their moves as compared to how they would react to indefinite deregulation, in order to arrive within the window. Footloose immigrants with relatively weak preferences across regions may demand entry into newly opened areas at a much higher level than they would if the deregulation was more general. This effect would, of course, lead a naive planner to overestimate the consequences of deregulation. The true effect would depend on the elasticities of substitution across regions, which may be unknowable.

The immigration example is not merely a thought experiment. In 2004, the EU expanded to include the so-called “A8” countries of Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia, and Slovenia. Accession nationals were formally granted the same rights of free immigration as nationals of extant members. As the accession approached there was widespread worry in the more-developed EU15 countries that they would experience a huge spike of immigration from new member states, with new immigrants competing

for jobs, depressing wages, and disrupting social cohesion. In response, the Treaty of Accession allowed EU15 members to impose “temporary” restrictions on worker immigration from the A8 countries for up to seven years after the accession. In the years immediately after accession only the UK, Ireland, and Sweden allowed open access to their labor markets, while the remaining A15 members maintained relatively strict work permit systems. A similar pattern held when Bulgaria and Romania (the “A2” group) were admitted to the union in 2006.

Prior to the opening an estimated 50,000 A8 and A2 nationals were residing in the UK, out of about 850,000 in the EU15 at large (Brücker, Alvarez-Plata and Siliverstovs 2003). Predictions of expected flows to the UK from the A8 ranged from 5,000 to 17,000 annually (Dustmann, Fabbri and Preston 2005). In reality, the immigrant flows were much larger than that. Even by the strictest definition — those who self-identify upon arrival that they intend to stay for more than a year — A8 immigration was 52,000 in 2004, 76,000 in 2005, and 92,000 in 2006 (Office for National Statistics 2006). Using estimates based on the Eurostat Labour Force Survey, Gligorov (2009) finds that net flow of A8 worker immigrants between 2004 and 2007 was just under 500,000.

One of the most cited explanations for the underestimate of immigration flows to the UK was not sufficiently accounting for the effects of the maintenance of immigration restrictions by the remaining 80-percent of the EU15 (Gilpin, Henty, Lemos, Portes and Bullen 2006). The traditional destinations for migrant workers from Eastern Europe, Germany and Austria, were closed off by the temporary continuance of immigration restriction. Instead of waiting for these countries to open up, the migrants instead came to the UK (and Ireland and Sweden, to a lesser degree). Not only is it hard for other Western European countries to learn much from the UK’s experiment, since they are not identically economically situated, but it’s even hard for the UK to learn much about what completely free immigration across the EU would mean for itself. The observed patterns are likely an overestimate of the effect the UK should expect from open borders, but the degree of overestimation will depend on how many of the migrants were crowded in by restriction elsewhere and how many legitimately preferred coming to the UK.

4.1 Optimal Policy with Regulation and Policy Experiments

Consider an augmentation of the base model presented in section 2 that expands the set of actions available to the planner in each period. In addition to regulating or deregulating, the planner may instead opt for a third option of performing an experiment. When performing the experiment, the planner continues to perform the primary regulatory function at cost d ,

and pays an additional cost of $c \geq 0$ to fund and monitor the experiment. The case of basic regulation with some costless information revelation is simply $c = 0$.

We consider the simplest signal structure from experimentation which captures the notion of imprecise information. An experiment can either be a success or a failure, depending on the underlying state and on chance. If the state is bad, the regulatory experiment will always be a failure. If the state is good, the regulatory experiment will succeed with probability α , and will fail with probability $(1 - \alpha)$.¹⁸

A planner observing a failed experiment can not determine whether this failure was due to a randomly failed experiment or a bad state of the world. Denote the updated beliefs from a failed experiment as $\hat{\epsilon}$. Then:

$$(8) \quad \hat{\epsilon} = \frac{\epsilon}{\epsilon + (1 - \epsilon)(1 - \alpha)}.$$

A planner observing a successful experiment will know she is in the good state for certain.

Let $E(\epsilon) \in \{0, 1\}$ represent the planner's experimentation strategy when she believes the state is bad with probability ϵ , where $E = 1$ indicates experimentation and $E = 0$ indicates no experimentation. Let $V(R, E|\epsilon)$ be the planner's value function playing regulation strategy R and experimentation strategy E with beliefs ϵ . Define $V^{**}(\epsilon)$ as the value function of a planner who chooses the maximizing regulation and experimentation regime, and let $\{R^{**}(\epsilon), E^{**}(\epsilon)\}$ be that maximizing strategy. Since experimentation yields strictly less information than deregulation and has an additive cost, deregulating and experimenting in the same period is never optimal.

Given maximization in all subsequent periods for any belief ϵ , the value for regulation, deregulation, and experimentation are respectively:

$$\begin{aligned} V(R = 1, E = 0|\epsilon) &= -d + \delta V^{**}(P(\epsilon)), \\ V(R = 0, E = 0|\epsilon) &= \epsilon[-1 + \delta V^{**}(P(1))] + (1 - \epsilon)[\delta V^{**}(P(0))], \\ V(R = 1, E = 1|\epsilon) &= -d - c + [\epsilon + (1 - \alpha)(1 - \epsilon)][\delta V^{**}(P(\hat{\epsilon}))] + (1 - \epsilon)\alpha[\delta V^{**}(P(0))]. \end{aligned}$$

As before, $V(R = 1, E = 0|\epsilon)$, $V(R = 0, E = 0|\epsilon)$, $V(R = 1, E = 1|\epsilon)$, and $V^{**}(\epsilon)$ are all continuous and weakly decreasing in ϵ . Further, since $1 > d > 0$, and $d + c \geq d$, deregulation is optimal at $\epsilon = 0$ and regulation is optimal at $\epsilon = 1$.

¹⁸The asymmetry here is for simplicity. If a success could sometimes occur with a bad state, the planner would update to some $\epsilon > \rho_{GB}$ and would either still deregulate immediately or proceed through the updating from that point, as appropriate. This alternative formulation would change none of the basic trade-offs identified and simply complicate the math in obvious ways.

Our solution strategy is similar to the base case in that we look for a cutoff belief ϵ^{**} such that the planner prefers experimentation to regulation when $\epsilon < \epsilon^{**}$. If this belief exists and is (1) greater than the cutoff point ϵ^* for which deregulation is better than regulation and (2) small enough that $P(\hat{\epsilon}^{**}) > \epsilon^{**}$, optimal policy calls for experimentation each time the planner's belief falls below ϵ^{**} and deregulation if this experimentation is a success. As $\hat{\epsilon} < 1$, a planner who is unsuccessful in experimentation will wait for a shorter amount of time before experimenting again. Thus optimal policy will typically be characterized by a long initial regulation period followed by cycles of experimentation and shorter regulatory spells.

Just as in the base model, the added cost of experimentation results in a larger set of parameters for which regulation is permanent, relative to the full-information baseline. If $\epsilon^{**} < \tilde{\epsilon}$, the planner's future value from experimentation is never high enough to justify the additional costs of being in the bad state. Letting ϵ^{**} converge to $\tilde{\epsilon}$ from above and continuing to assume $\epsilon^* < \epsilon^{**}$, regulation is permanent if $d \leq \rho_{GB}\tau'$, where

$$(9) \quad \tau' \equiv 1 + \left(\frac{c}{\alpha}\right) \left(\frac{1}{\kappa(1 - \tilde{\epsilon})}\right) > 1$$

and where $\kappa \equiv \frac{\delta\rho_{GB}}{1-\delta+\delta\rho_{GB}}$ denotes the expected cost of the first bad state discounted one period into the future.¹⁹ As is evident in the last term on the right hand side, permanent regulation is mitigated if the cost of experimentation — which has precision bounded at $\alpha(1 - \tilde{\epsilon})$ — is low relative to the value of information, which is bounded at $\frac{\rho_{GB}}{\kappa}$. As $\tilde{\epsilon}$ is the most optimistic belief that can occur under regulation, the value of information is maximal at this point. Thus if experimentation does not have a positive net present value at the steady state, it never will.

On the other hand, if $\epsilon^{**} < \epsilon^*$, the planner's optimal policy involves only deregulation. In this case experimentation will never be used, and optimal policy is identical to that found in section 2.

Finally, there is also a hybrid case which can occur if $P(\hat{\epsilon}^{**}) < \epsilon^{**}$ and $\epsilon^* \leq \epsilon^{**}$. In this case a planner may find it in her interest to continuously experiment at ϵ^{**} and below, but eventually deregulate if her beliefs fall below a secondary threshold. This case only occurs if α is extremely low or d is very high. We view this case as most relevant when thinking about situations in which planners may be privy to external information in each period and return to the special case in section 4.2.

Proposition 4 summarizes the optimal policy with experimentation:

¹⁹ κ is derived in the appendix as part of Proposition 3.

Proposition 4 *Assume the planner has access to deregulatory experiments, at cost c , which will succeed with probability α in the good state and probability 0 in the bad state. Let τ be defined as in Proposition 3 and let τ' be defined as above. Once regulation is applied for the first time, the planner's optimal policy falls into one of the following cases:*

1. *If $d \leq \rho_{GB} * \min\{\tau, \tau'\}$, the planner regulates every period and never experiments.*
2. *If $d > \rho_{BB}$, the planner never regulates or experiments, even after a bad state.*
3. *If $\rho_{GB} * \min\{\tau, \tau'\} < d \leq \rho_{BB}$, the planner regulates for k^{**} periods and then either deregulates or experiments. If the optimal choice is to deregulate, then $k^{**} = k^*$ from Proposition 3. If the optimal choice is to experiment, then k^{**} is the first k such that $P^{k+1} \leq \epsilon^{**} \leq P^k(1)$ and ϵ^{**} is the solution to the implicit function:*

$$(10) \quad c + \delta V^{**}(P(\epsilon)) - [\epsilon + (1 - \alpha)(1 - \epsilon)][\delta V^{**}(P(\hat{\epsilon}))] - (1 - \epsilon)\alpha[\delta V^{**}(P(0))] = 0$$

As with the case without experimentation, the range of costs for which regulatory fog leads to permanent persistence with experimentation is decreasing in δ . In the case where $\epsilon^* < \epsilon^{**}$ and experiments is used, the increase of persistence is due to a decrease in κ as δ increases. In the other two cases, the derivation is exactly the same as our baseline model. In all cases the intuition is straightforward: the results of an experiment is only valuable next period, so more patient regulators are more likely to use experiments.

Figure 5 shows the regions for which experimentation and deregulation is used for different regulation costs, d , and experimentation costs c in the case of highly informative experiments ($\alpha = .5$), moderate discounting ($\delta = .9$) and highly persistent states ($\rho_{GB} = \rho_{BG} = .05$). For low regulation costs (below $\tau\rho_{GB}$), deregulation is never used on its own, but only after a successful experiment, and the optimal policy decision is between permanent regulation and experimentation. When experimentation is costless, it will be used as long as deregulation is ever attractive ($d > \rho_{GB}$). As regulation becomes more expensive, the value of information increases leading to a greater value for experimentation and a concomitant increase in the acceptable costs.

For moderate to high regulation costs ($d > \rho_{GB}\tau$), a planner always deregulates eventually, and thus her decision is between implementing a strict policy of regulation and deregulation cycles or a policy which also includes experimentation. As d increases, the relative cost to deregulating declines and thus deregulation becomes strictly more attractive. As d approaches ρ_{BB} , the rush to deregulate leads to the abandonment of experimentation. The intuition here is that if the planner plans to deregulate next period, even after a failed

experiment, there is no reason to pay for an experiment.²⁰

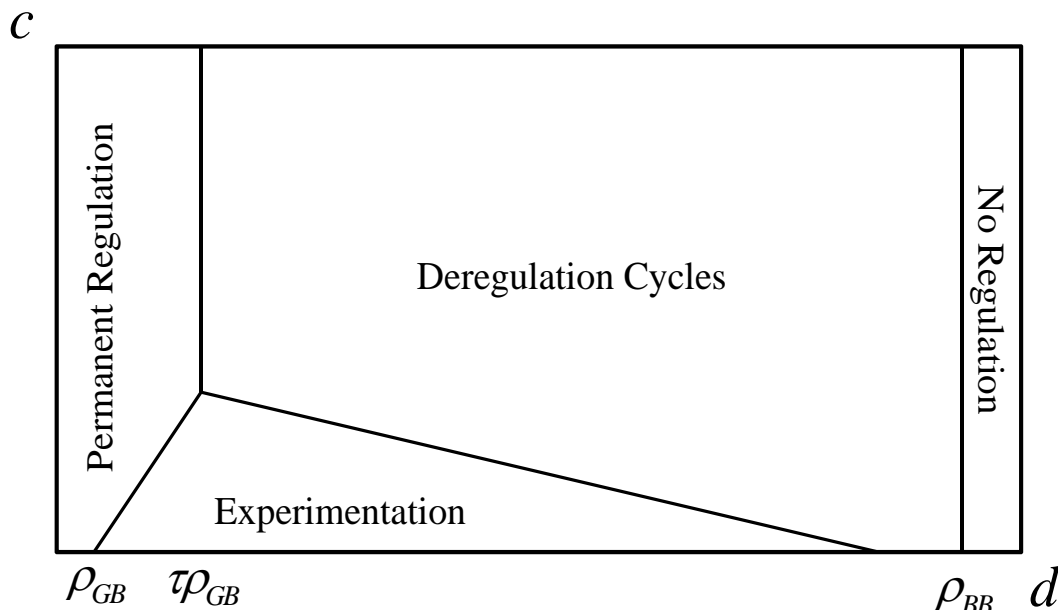


Figure 5: Experiments vs. Deregulation ($\rho_{GB} = \rho_{BG} = .05, \delta = .9, \alpha = .5$)

A special case of our model with experimentation that is of particular importance is the case where information arrival is exogenous, i.e., where the cost of information is zero. This variant of the model is likely a good representation of what occurred at the CAB as well as being a more realistic model of federalist regulation.

The presence of exogenous information has two effects in our model that combine to make the overall effect of exogenous information on policy persistence ambiguous. First, the presence of a non-zero probability of information arrival eliminates permanent persistence since sooner or later good news will arrive. Second, the value of regulating has increased, since there is now some chance of observing the state despite the regulation. This secondary effect lowers the relative value of experimentation and causes the length of a regulatory cycle to increase. Depending on the parameters of the model, exogenous information may crowd out experimentation and lead to quasi-permanent regulation— i.e., regulation that stays in place until good information is exogenously supplied. The manner in which deregulation cascaded through a plethora of network industries within a couple decades in the 1970s, suggests exactly this sort of phenomenon.

²⁰For very pessimistic beliefs, the update after a failed experiment is actually more optimistic than the prior belief, since the natural progression of the Markov process is quite large for beliefs far from the steady state. For example, it's easy to check that $P(\hat{\rho}_{BB}) < \rho_{BB}$ for any α .

4.2 The Value of Experimentation

Obviously, the ability to experiment has no value to the planner when it is never used in equilibrium, and has some positive value when used. When used, its value will depend in intuitive ways on the cost of running the experiment and the precision of information that it uncovers. While the cost of experimentation acts linearly on the value of experimentation, the precision of information does not. As $\frac{1}{\alpha}$ is multiplicative, experiments with low precision have very limited value to the policy maker. In these cases, the planner finds it in her interest to never use experimentation, or to use it only in conjunction with periodic deregulation.

Figure 6 shows the range of parameters for which experimentation and simple deregulation are preferred in $\{\alpha, c\}$ space. As the value of deregulation is a constant in this space, there exists an iso-efficiency “indifference” curve along which the value of the optimal strategy using experimentation is exactly equal to the optimal strategy without it. As the value of experimentation is increasing in α and decreasing in c , policies which are to the southeast of this curve are preferred to policies consisting only of deregulation. One way to interpret this curve, $c(\alpha)$, is as the planner’s “willingness to pay” for an experiment of a certain precision. It follows that any experiment falling below the line would net him a surplus at least proportional to the distance below the line.²¹ This surplus is the (net) value of experimentation.

When α is small, recall that there may be cases in which $P(\hat{\epsilon}^{**}) < \epsilon^{**}$ and thus the planner’s beliefs are improving even after a failed experiment. In these cases, optimal regulation may call for both experimentation and eventual deregulation. As can be seen in Figure 6, the region for which this occurs is for α and c very small. This case is likely to include situations such as the CAB example where the planner may learn about the state of nature from a low probability exogenous source. As $P(\hat{\epsilon})$ is always greater than $P(\epsilon)$, the time between deregulation experiments is increasing in the likelihood of external information. Thus, the possibility for external signals can actually increase the length of the regulatory spell even though it unambiguously improves welfare.

The observation that the optimal experimentation frontier can be expressed as an iso-efficiency “indifference” curve provides a method by which alternative policies can be evaluated. Consider a collection of N experiments, indexed by $i = 1, 2, \dots, N$, which are available for the planner to choose among. Each experiment consists of a (c_i, α_i) combination, and we will assume that the planner must choose one to use and stick with that choice whenever he decides to experiment. Then this framework provides a way of analyzing the planner’s

²¹It would be exactly proportional if the experimenting strategy was unchanged by the reduced cost, but the optimizing planner may also decide to start experimenting more often, further improving his payoff.

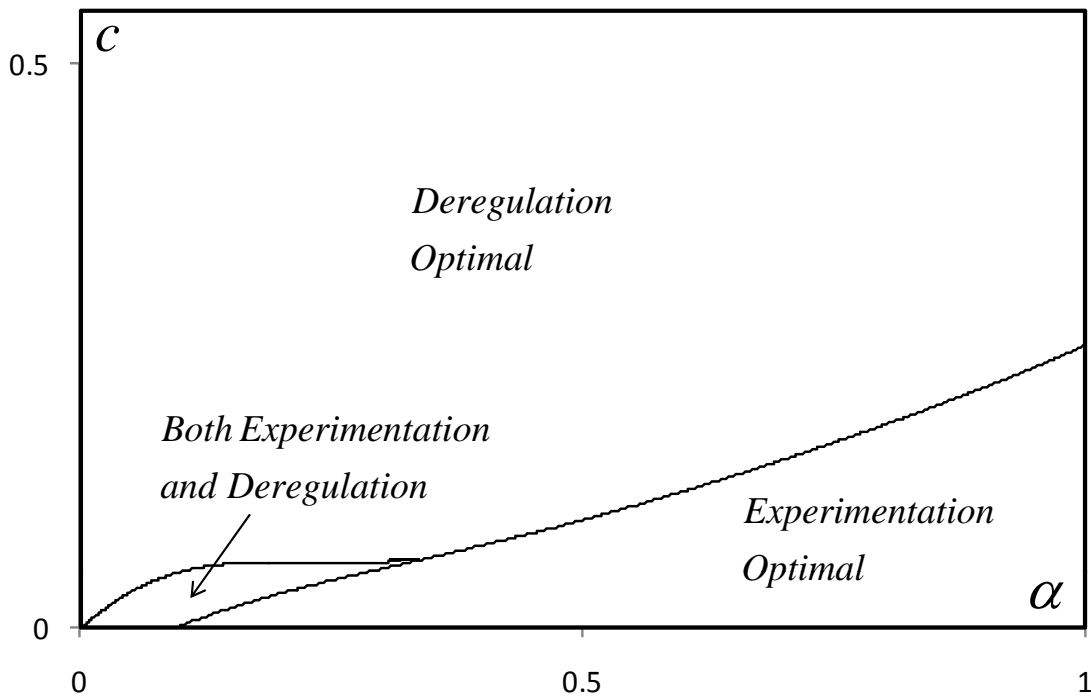


Figure 6: Experiments vs. Deregulation ($\rho_{GB} = \rho_{BG} = .05, \delta = .9, d = .5$)

optimal choice among these experiments. The frontier in Figure 6 is merely one (particularly salient) indifference curve for experiments. For any experiment (c_i, α_i) below that frontier we could identify a similar increasing curve $c_i(\alpha)$ which includes that experiment ($c_i(\alpha_i) = c_i$) and for which the planner is indifferent among all the experiments on the curve. Optimal choice for the planner, then, simply amounts to choosing the experiment on the lowest indifference curve. Since $(0, 0)$ is always in the set of experimental options, if all other options are above the curve in Figure 6, the optimal choice is simply to never experiment. If the feasible set of experiments is convex, the familiar tangency condition for indifference curves and the budget frontier will characterize the optimal choice. Of course, all the natural statics would follow from this characterization: more precise experiments should be preferred if the marginal cost of precision falls (budget curve gets less convex) and more precise experiments should be chosen as the marginal value of precision increases.²²

²²While we have concentrated our analysis on the case of temporary experiments, in order to build off the analysis in the preceding sections, a similar exercise could be done in order to compare various sorts of regulation, which differ with respect to how much information they let through. Imagine, for example, two methods of regulating. The first is exactly like the regulation described above. It costs d_1 to implement but shuts down all information. The other costs $d_2 > d_1$ to implement but reveals the good state as good with some probability, $\alpha > 0$. The difference between this way of posing the problem and the way we describe experiments is that the choice over regulatory regimes would be made ex-ante, and the higher price of the informative regulation would need to be paid every period that regulation is imposed, instead of a temporary premium for a temporary experiment. The most natural way of modeling the choice would depend on the

As the cost of deregulation is a function of d , the location of the frontier between deregulation and experimentation is also a function of the cost of regulation. As we saw from Proposition 4, experiments are never used for $d < \rho_{GB}$ or $d > \rho_{BB}$, so the planner would never be willing to pay for information in those cases (i.e., $c(\alpha)$ is flat at zero). Furthermore, as was outlined in Remark 1, the costs of regulatory fog are most severe for intermediate costs of regulation due to the high frequency of failed policy experiments and the burden of regulation. It is precisely in these states that the willingness to pay for information is highest, and thus $c(\alpha)$ is highest for intermediate d . As $c(0) = 0$ for all d , the iso-efficiency curves for experimentation will become steeper as d approaches the value which maximizes the cost of regulatory fog. It follows that the value of an additional unit of precision is always highest when d is closest to the point which maximizes the losses due to regulatory fog in the baseline case.

Our extended model suggests that small-scale experimentation is not a panacea for regulatory fog. As in the CAB example, the potential for external information may actually increase the expected length of regulatory spells since regulators may prefer to wait for external information rather than risk a regulatory disaster. Further, as in the immigration example at the beginning of the section, experimentation may be very costly and generate limited information when the policy is meant to mitigate moral hazard in the target population.

5 Conclusions

Models of regulatory persistence are typically based on the role that agency and lobbying play in influencing final policy or on some technical fixed cost of changing policy. We argue that in many environments, regulation generates the seeds of its own persistence by altering the information observable about the environment — a phenomenon we refer to as regulatory fog. Under a stark policy environment of regulation and deregulation and in a broader environment where experimentation is also allowed, we find that the effects of regulatory fog can be quite severe. Regulatory fog can lead to permanent regulation for a broad range of parameters, particularly by impatient planners. For most reasonable parameter values, fog delays deregulation and causes the economy to stay in the regulated state more often than the underlying environment warrants alone. Finally, fog can lead to deregulatory disasters

technology at hand. If switching among regulatory regimes is very costly, this second model may be more appropriate. Nevertheless, the results are quite similar using this alternative approach. We would again end up with indifference curves in the (d, α) space with roughly the same shape as those appearing in Figure 6, and the tradeoffs that guide optimal choice amongst regulatory regimes would be quite similar to those discussed here.

which can greatly diminish overall social welfare.

While we have intentionally chosen the simplest of information structures to analyze, the intuition from our analysis extends to much more complicated information structures. Just as in our model, at each point in time the planner faces a tradeoff in the cost and benefit of each potential policy and the informativeness of the policy for future decision making. More informative policies generate a higher future expected value, but also lead to much faster responses to changes in the underlying environment. The slow and imperfect update process associated with less informative policies often lead to delays as optimal planners prefer to decrease risk and increase the value of information. Our analysis demonstrates that this delay can lead to increased persistence of policies which are uninformative, a result that is quite robust to alternative information structures and policies.

Although we have chosen to explore regulatory fog in an environment with a perfectly public-interested planner, the information and political economy channels are quite complementary. In an interest group model such as Coate and Morris (1999), information asymmetries between regulated firms and consumers are likely to generate significant pressure from regulated firms who are enjoying the protections of a regulated monopoly, but limited pull by consumers who are uncertain as to the final outcome of deregulation. Likewise, in an environment with politically charged regulation, partisan policymakers may develop policies which deliberately eliminate information in order to limit the ability of competing parties to overturn legislation in the future.

While our analysis assumes a centralized planner, decentralization is of limited use when separated districts are symmetric and competitive. As pointed out by Rose-Ackerman (1980) and generalized by Strumpf (2002), the potential policy experiments in other districts provides incentives for policy makers to delay their own deregulatory policies and can, in many cases, actually lead to more regulatory persistence.²³ Further, just like in the experimentation example, spill-overs from one district to another are likely to reduce the informativeness of experimentation and may ultimately make unilateral policy decisions fail.

Our research opens both theoretical and empirical avenues for future work. On the theoretical frontier, our current paper has restricted attention to the case where the agents do not have incentives to strategically respond to regulation. In a world where agents are long-lived and can influence information, it would be interesting to see to what extent regulatory fog affects strategic interaction. When agents have a preference for being regulated (as in a protectionist regulated monopoly setting), models of signal-jamming suggest that

²³This theme is echoed in the social learning literature where social learning leads to strategic delay in experimentation. See, for example, Gale (1996); Moscarini, Ottaviani and Smith (1998); Veldkamp (2005); and Peck and Yang (2010).

regulatory fog would be exacerbated. We have also restricted attention to the case where the regulator is benevolent and does not have control over the informativeness of policy. In a partisan landscape, we believe that bureaucrats may intentionally eliminate information from policies in order to extend their policies beyond their own tenure.

From an empirical perspective, our paper generates a number of testable hypothesis which we believe could be empirically explored. From a political economy standpoint, our model would predict that variation in term lengths and term limits should have an impact on deregulation policy. Regulators with short terms in office or who are concerned with reelection should be less likely to experiment with deregulation relative to their long-lived counterparts. Extensions to our model would also suggest that policy persistence may be exacerbated by federalism and that partisan politicians may favor policies that generate regulatory fog if they are likely to be voted out of office.

Looking across industries, our model would predict that as the cost of regulation increases and the cost of regulatory disasters decrease, both experimentation and the frequency of regulatory disasters should increase. Our model further predicts that regulation that either proscribes or deters actions are likely to be the most persistent. These type of regulation include the regulation of network industries, the regulation of illegal activities, pollution enforcement, or regulation put in place to eliminate moral hazard. Our model also predicts deregulatory cascades where regulators learn from the successful deregulation of similar industries.

Finally, although this analysis has focused on regulation, we believe regulatory fog is a general phenomenon which affects a wide variety of economic environments. Many economic institutions such as monitoring, certification, intermediation, and organizational structures are designed to alter the actions of heterogeneous agents which, in the process, affects the dynamic information generated. These dynamic effects are likely to influence both the long-term institutions which persist and the overall structure of markets and organizations.

References

- Aidt, Toke S. and Jayasri Dutta**, “Policy Myopia and Economic Growth,” *European Journal of Political Economy*, 2007, 23, 734–53.
- Bailey, Elizabeth**, “Deregulation and Regulatory Reform: U.S. Domestic and International Air Transportation Policy During 1977-79,” in “Regulated Industries and Public Enterprises: European and United States Perspective,” Lexington, 1980.

- Barrett, Sean D.**, “Exporting Deregulation — Alfred Kahn and the Celtic Tiger,” *Review of Network Economics*, 2008, 7, 573–602.
- Berentsen, Aleksander, Esther Bruegger, and Simon Loertscher**, “Learning, public good provision, and the information trap,” *Journal of Public Economics*, 2008, 92, 998–1010.
- Bergemann, Dirk and Juuso Välimäki**, “Bandit Problems,” *Cowles Foundation Discussion Paper No. 1551*, 2006.
- Boyer, Kenneth D.**, “The Costs of Price Regulation: Lessons from Railroad Deregulation,” *The RAND Journal of Economics*, 1987, 18, 408–416.
- Brainard, S. Lael and Thierry Verdier**, “Lobbying and Adjustment in Declining Industries,” *European Economic Review*, 1994, 38, 586–95.
- Brücker, Herbert, Patricia Alvarez-Plata, and Boriss Siliverstovs**, “Potential Migration from Central and Eastern Europe into the EU-15 An Update,” Technical Report, European Commission, DG Employment and Social Affairs 2003.
- Coate, Stephen and Stephen Morris**, “Policy Persistence,” *American Economic Review*, 1999, 89, 1327–1336.
- Derthick, Martha and Paul J. Quirk**, *The Politics of Deregulation*, The Brookings Institute, 1985.
- Dustmann, Christian, Frecesca Fabbri, and Ian Preston**, “The Impact of Immigration on the British Labour Market,” *Economic Journal*, 2005, 115, 324–341.
- Fernandez, Raquel and Dani Rodrik**, “Resistance to Reform: Status Quo Bias in the Persistence of Individual-Specific Uncertainty,” *American Economic Review*, 1991, 81, 1146–55.
- Friedrich, Silke**, “Policy persistence and rent extraction,” *Public Choice*, 2012, pp. 1–18. 10.1007/s11127-012-9945-9.
- Gale, Douglas**, “Delay and Cycles,” *Review of Economic Studies*, 1996, 63, 169–198.
- Gilpin, Nicola, Matthew Henty, Sara Lemos, Jonathan Portes, and Chris Bullen**, “The Impact of Free Movement of Workers from Central and Eastern Europe on the UK Labour Market,” 2006. Working Paper No. 29, Department of Work and Pensions, UK.
- Gligorov, Vladimir**, “Mobility and Transition in Integrating Europe,” 2009. UNDP Human Development Research Paper 2009/15.

- Grossman, Gene M. and Elhanan Helpman**, “Protection for Sale,” *American Economic Review*, 1994, *84*, 833–50.
- Hirshleifer, Jack and John G. Riley**, “The Analytics of Uncertainty and Information — An Expository Survey,” *Journal of Economic Literature*, 1979, *17*, 1375–1421.
- and — , *The Analytics of Uncertainty and Information*, Cambridge University Press, 1992.
- Ichino, Andrea and Gerd Muehlheusser**, “How often should you open the door? Optimal monitoring to screen heterogeneous agents,” *Journal of Economic Behavior & Organization*, 2008, *67*, 820–831.
- Jehiel, Philippe and Andrew F. Newman**, “Loopholes: Social Learning and the Evolution of Contract Form,” 2011. mimeo.
- Jordan, William**, *Air Regulation in America: Effects and Imperfections*, The Johns Hopkins Press, 1970.
- Joskow, Paul**, “Markets for Power in the United States: An Interim Assesment,” *The Energy Journal*, 2006, *1*, 1–36.
- , “Lessons Learned from Electricity Market Liberalization,” *The Energy Journal*, 2008, *Special Issue on the Future of Electricity*, 9–42.
- Leaver, Clare**, “Bureaucratic Minimal Squawk Behavior: Theory and Evidence from Regulatory Agencies,” *American Economic Review*, 2009, *99*, 572–607.
- Levine, Michael E.**, “Is Regulation Necessary? California Air Transportation and National Regulatory Policy,” *Yale Law Journal*, 1965, *74*, 1416–47.
- McLennan, Andy**, “Price Dispersion and Incomplete Learning in the Long Run,” *Journal of Economic Dynamics and Control*, 1984, *7*, 331–347.
- Moore, Thomas Gale**, “Rail and Truck Reform: The Record so Far,” *Regulation*, 1988, *6*, 57–62.
- Morrison, Steven**, “Airline Service: The Evolution of Competition since Deregulation,” in Victor J. Tremblay and Carol Horton Tremblay, eds., *Industry and Firm Studies*, M.E. Sharpe, 2007.
- and **Clifford Winston**, *The Evolution of the Airline Industry*, The Brookings Institute, 1995.
- Moscarini, Giuseppe, Marco Ottaviani, and Lones Smith**, “Social Learning in a Changing World,” *Economic Theory*, 1998, *11*, 657–665.

- Panzar, John C.**, “Regulation, Deregulation, and Economic Efficiency: The Case of the CAB,” *American Economic Review*, 1980, *70*, 311–15.
- Peck, James and Huanxing Yang**, “Investment Cycles, Strategic Delay, and Self-Reversing Cascades,” *mimeo*, 2010.
- Peltzman, Sam and Clifford Winston**, *Deregulation of Network Industries: What’s Next?*, The Brookings Institution, 2000.
- , **Michael Levine, and Roger G. Noll**, “The Economic Theory of Regulation after a Decade of Deregulation,” *Brookings Papers on Economic Activity: Microeconomics*, 1989, *1989*, 1–59.
- Piketty, Thomas and Emmanuel Saez**, “How Progressive is the U.S. Federal Tax System? A Historical and International Perspective,” *Journal of Economic Perspectives*, 2007, *21*, 3–24.
- Pindyck, Robert S.**, “Irreversibilities and the Timing of Environmental Policy,” *Resource and Energy Economics*, 2000, *22*, 233–259.
- Radner, Roy and Joseph E. Stiglitz**, “A Nonconcavity in the Value of Information,” in R.E. Kihlstrom, ed., *Bayesian Models in Economic Theory*, Elsevier Science Publishers, 1984.
- Office for National Statistics**, “International Migration: Migrants Entering or Leaving the United Kingdom, England and Wales,” Technical Report, Office for National Statistics 2006.
- Robbins, Herbert**, “Some Aspects of the Sequential Design of Experiments,” *Bulletin of the American Mathematical Society*, 1952, *55*, 527–535.
- Rose-Ackerman, Susan**, “Risk Taking and Reelection: Does Federalism Promote Innovation,” *Journal of Legal Studies*, 1980, *9*, 593–616.
- Rose, Nancy L.**, “The Incidence of Regulatory Rents in the Motor Carrier Industry,” *The RAND Journal of Economics*, 1985, *16*, 299–316.
- Rubinovitz, Robert N.**, “Market Power and Price Increases for Basic Cable Service since Deregulation,” *The RAND Journal of Economics*, 1993, *24*, 1–18.
- Smith, Janet K.**, “Production of Licensing Legislation: An Economic Analysis of Interstate Differences,” *Journal of Legal Studies*, 1982, *11*, 117–138.
- Stigler, George J.**, “The Theory of Economic Regulation,” *Bell Journal of Economics*, 1971, *2*, 3–21.

Strumpf, Koleman, “Does Government Decentralization Increase Policy Innovation?,” *Journal of Public Economic Theory*, 2002, 4, 207–241.

Veldkamp, Laura L., “Slow Boom, Sudden Crash,” *Journal of Economic Theory*, 2005, 124, 230–257.

Winston, Clifford, “U.S. Industry Adjustment to Economic Deregulation,” *Journal of Economic Perspectives*, 1998, 12, 89–110.

—, **Thomas M. Corsi, Curtis M. Grimm, and Carol A. Evans**, *The Economic Effects of Surface Freight Deregulation*, The Brookings Institution, 1990.

Zhao, Jinhua and Catherine L. Kling, “Policy persistence in environmental regulation,” *Resource and Energy Economics*, 2003, 25, 255–268.

6 Appendix

6.1 Proofs from Main Text

PROPOSITION 1:

Proof. We begin by defining a Markov transition matrix Q with the following four states: s_1 is the state in which the last two periods were Good; s_2 is the state in which the last period was good and the previous period was bad; s_3 is the state in which the last two periods were bad; and s_4 is a state in which last period was bad and the previous period was good. As the selected policy has no influence over the transmission between these states, the transition matrix between states for any potential policy is given as follows:

$$(11) \quad Q = \begin{pmatrix} \rho_{GG} & 0 & 0 & \rho_{GB} \\ \rho_{GG} & 0 & 0 & \rho_{GB} \\ 0 & \rho_{BG} & \rho_{BB} & 0 \\ 0 & \rho_{BG} & \rho_{BB} & 0 \end{pmatrix},$$

A policy under full information is defined as a cost vector corresponding to each of the potential states. For a policy that always regulates, the cost vector $c_{FullReg}$ is simply $(-d \ -d \ -d \ -d)$. Likewise, a policy which never regulates has a cost vector $c_{NeverReg}$ given by $(0 \ 0 \ -1 \ -1)$. Finally, a policy which regulates after a bad state and has no regulation after a good state, c_{cycle} yields $(0 \ -d \ -d \ -1)$.

For any initial probability vector v , the expected cost of a policy c is given by:

$$(12) \quad \sum_{i=1}^{\infty} \delta^{(i-1)} v Q^i c.$$

This is simply the cost of each potential policy in each potential state weighted by the probability of arriving in that state. As the Markov matrix is the same for each policy, the potential trade off of each policy can be seen by subtracting the cost functions. The cyclical policy is superior to full regulation when:

$$(13) \quad \sum_{i=1}^{\infty} \delta^{(i-1)} v Q^i [c_{cycle} - c_{FullReg}] \geq 0.$$

Noting that $[c_{cycle} - c_{FullReg}] = (d \ 0 \ 0 \ -(1-d))$ and that states s_1 and s_4 can only be entered into by s_1 and s_2 , where the state of last period was good, it follows that cyclical policy is superior to full regulation iff:

$$(14) \quad d\rho_{GG} - (1-d)\rho_{GB} > 0 \rightarrow d > \rho_{GB}$$

Repeating the exercise for a policy which never regulates, $[c_{cycle} - c_{FullReg}] = (0 \ -d \ 1-d \ 0)$. As states s_2 and s_3 can only be entered into by s_3 and s_4 , where the state of last period was bad, it follows that cyclical policy is superior to no regulation iff:

$$(15) \quad -d\rho_{BG} + (1-d)\rho_{BB} > 0 \rightarrow \rho_{BB} > d$$

As the matrix Q and the cost vector c captures both the underlying transition matrix and the cost of (i) being regulated in a good state and (ii) being unregulated in a bad state, it captures all payoff relevant information of a policy. Any other candidate policy which deviates from one of the three above must do so in a state with identical payoff-relevant properties as one of the states of Q . As the chosen policy will be optimal in this state, the candidate policy must be strictly dominated.

Under the cyclical policy, the probability of continuing regulation is exactly the probability of staying in the bad state, ρ_{BB} . So the probability of having a spell of length t is given by $\rho_{BB}^{t-1}(1-\rho_{BB})$. This is exactly the pdf of a random variable with a geometric distribution with parameter ρ_{BB} , which has a mean length of $1/(1-\rho_{BB})$. Finally, the time spent in regulation is simply the stationary distribution of states s_1 and s_2 of Q . As Q is regular, the stationary distribution is given by the vector π s.t. $\pi Q = \pi$. Solving this completes the proof. ■

PROPOSITION 2:

Proof. Assume that the planner has some optimal strategy $R^*(\epsilon)$ which induces a value function $V^*(\epsilon)$. For any ϵ define

$$(16) \quad G(\epsilon) = V(R = 1|\epsilon) - V(R = 0|\epsilon).$$

$V(R|\epsilon)$ is continuous in ϵ and thus G is continuous. Since $G(0) < 0$, $G(1) > 1$, and G is continuous, there is some ϵ^* for which $G(\epsilon^*) = 0$. For the Proposition it would suffice to show that this ϵ^* is unique. In fact, we show that $G(\epsilon)$ is increasing, a stronger claim. Replacing for (3) and (4) in equation (16),

$$G(\epsilon) = -d + \delta V^*(P(\epsilon)) - \epsilon[-1 + \delta V_B] - (1 - \epsilon)\delta V_G.$$

Replacing in turn for $V^*(P(\epsilon))$, this becomes

$$(17) \quad \max \left\{ \begin{array}{l} -d + \delta[(-d + \delta V^*(P^2(\epsilon)))] - \epsilon[-1 + \delta V_B] - (1 - \epsilon)\delta V_G, \\ -d + \delta[P(\epsilon)[-1 + \delta V_B] + (1 - P(\epsilon))\delta V_G] - \epsilon[-1 + \delta V_B] - (1 - \epsilon)\delta V_G \end{array} \right\},$$

where the first constituent of the maximand is the difference in returns when choosing to regulate next period after regulating this period versus not regulating this period, and the second constituent is the return to not regulating next period after regulating this period versus not regulating this period. More generally, define $G^k(\epsilon)$ as the difference between the return for regulating for k periods and then following optimal strategies from then on and simply not regulating this period. I.e.,

$$G^k(\epsilon) = -d \sum_{j=0}^k \delta^{j-1} + \delta^k \{P^k(\epsilon)[-1 + \delta V_B] + (1 - P^k(\epsilon))\delta V_G\} - \epsilon[-1 + \delta V_B] - (1 - \epsilon)\delta V_G$$

Then for all k , $G^k(\epsilon)$ is differentiable and increasing. Furthermore

$$\lim_{k \rightarrow \infty} G^k(\epsilon) = -\frac{d}{1 - \delta} - \epsilon[-1 + \delta V_B] - (1 - \epsilon)\delta V_G,$$

which is also increasing in ϵ . Finally, note that $G(\epsilon) = \max_k G^k(\epsilon)$, since by assumption that regulatory will act optimally in subsequent periods. Since $G(\epsilon)$ and $G^k(\epsilon)$ are all continuous, it follows that $G(\epsilon)$ must also be increasing. Therefore, there is a unique ϵ^* where $G(\epsilon) \geq 0$ if and only if $\epsilon \geq \epsilon^*$, and that ϵ^* will, therefore, satisfy the requirements of the Proposition. ■

COROLLARY 1:

Proof. We'll prove this using the implicit function theorem on $G(\epsilon)$. From the proof of Proposition 2, $G'(\epsilon) > 0$, and so it follows directly from the implicit function theorem that

$$(18) \quad \text{sgn}\left(\frac{\partial \epsilon^*}{\partial d}\right) = \text{sgn}\left(-\frac{\partial G(\epsilon, d)}{\partial d}\right)$$

$\frac{\partial G(\epsilon, d)}{\partial d}$ may vary depending on whether ϵ^* is below the steady state $\tilde{\epsilon}$ or above it. The easier case occurs if $\epsilon^* \leq \tilde{\epsilon}$. Here, $P(\epsilon^*) \geq \epsilon^*$, and so $V(P(\epsilon^*)) = V_B = \frac{-d}{1-\delta}$. In this case,

$$(19) \quad G(\epsilon, d) = -d - \delta\left(\frac{d}{1-\delta}\right) - \epsilon\left(-1 - \delta\frac{d}{1-\delta}\right) - (1-\epsilon)\delta\left[\frac{-\kappa}{\delta} - \kappa\left(\frac{d}{1-\delta}\right)\right].$$

and so

$$(20) \quad \frac{\partial G(\epsilon, d)}{\partial d} = -1 + \frac{\delta}{1-\delta}\left[-1 + \epsilon + (1-\epsilon)\kappa\right] < 0.$$

If $\epsilon^* > \tilde{\epsilon}$, then $P(\epsilon^*) < \epsilon^*$, and so $V(P(\epsilon^*)) = P(\epsilon^*)(-1 + \delta V_B) + (1 - P(\epsilon^*))\delta V_G$. Noting that $V_G = \kappa[-\frac{1}{\delta} + V_B]$, and thus $\frac{\partial V_G}{\partial d} = \kappa\frac{\partial V_B}{\partial d}$,

$$(21) \quad \frac{\partial G(\epsilon^*, d)}{\partial d} = -1 + \delta\frac{\partial V_B}{\partial d}\left[\delta[P(\epsilon^*) + (1 - P(\epsilon^*))\kappa] - [\epsilon^* + (1 - \epsilon^*)\kappa]\right].$$

Since regulation cannot be used more than once per-period, $\frac{\partial V_B}{\partial d} > -\frac{1}{1-\delta}$. Furthermore, since $P(\epsilon^*) < \epsilon^*$, $P(\epsilon^*) + (1 - P(\epsilon^*))\kappa < \epsilon^* + (1 - \epsilon^*)\kappa$, and so

$$\frac{\partial G(\epsilon^*, d)}{\partial d} < -1 + \delta\left(\frac{-1}{1-\delta}\right)(\epsilon^* + (1 - \epsilon^*)\kappa)(\delta - 1) = -1 + \delta(\epsilon^* + (1 - \epsilon^*)\kappa) < 0.$$

■

COROLLARY 2:

Proof. This follows directly from Corollary 1. The update operator $P()$ is unaffected by d , and ϵ^* increases. Thus k^* must be weakly decreasing in d . ■

PROPOSITION 3:

Proof. In deriving this strategy, there are two cases to consider. First, assume that $\epsilon^* \leq \rho_{GB}$. Here, even after observing the good state, the planner will want to regulate. Since the planner takes the same action in the good and the bad states, $V_G = V_B = V^*(P(\epsilon))$ and thus the value of information is zero. Thus $G(\epsilon^*) = 0$ when $\epsilon^* = d$, so this case will occur only if $d \leq \rho_{GB}$, just like in the Full-Information benchmark.

Looking at the more interesting case, assume that $\epsilon^* > \rho_{GB}$, so regulation will not be imposed in the period immediately after the good state is observed. In this case equilibrium regulation has the following simple structure. After observing the bad state, the planner will regulate for k^* periods (perhaps infinite) and deregulate in the $(k^* + 1)$ period to see if the state has changed. If, upon sampling, she observes the bad state, she updates her posterior to $P(1)$ at the start of the next period and begins the regulation phase again. If, on the other hand, she finds herself in the good state, she does not regulate again until she experiences the bad state.

For $\epsilon^* > P(1) = \rho_{BB}$ this strategy means the planner actually never imposes regulation. As the value of information in this case is zero, the no regulation criterion is the same as the full information model, with no regulation imposed when $d > \rho_{BB}$.

For $\epsilon^* \in (\rho_{GB}, \rho_{BB}]$, a planner who arrives in the bad state will impose regulation and lift it every $k^* + 1$ periods to see if the state has changed. This region is characterized by potentially long (or infinite) periods of regulation, punctuated by deregulation at fixed intervals. If $\epsilon^* \leq \tilde{\epsilon}$, the planner's beliefs will converge to the stationary state which is above the cutoff necessary for deregulation. The planner's future value from deregulating is not high enough to justify the potential risk of being in the bad state.

To differentiate between the permanently persistent regulation case and the regulatory cycles case, it suffices to find the parameter values for which ϵ^* converges to $\tilde{\epsilon}$ from above. In the region of mixed regulation V_G is related to V_B by the potential transition from the good to the bad state. Let κ denote the expected cost of the first bad state discounted one period into the future:

$$(22) \quad \kappa = \sum_{t=0}^{\infty} \delta^t (1 - \rho_{GB})^t \delta \rho_{GB} = \frac{\delta \rho_{GB}}{1 - \delta + \delta \rho_{GB}}.$$

The expected value of the period following the good state is given by

$$(23) \quad V_G = V^*(P(0)) = \rho_{GB}[-1 + \delta V_B] + \rho_{GG} \delta V_G = \kappa[-1/\delta + V_B],$$

where the first term is the cost of being caught in the bad state without regulation and the second term is the future valuation of being in the bad state with certainty.

As ϵ^* converges to $\tilde{\epsilon}$ from above, $k^* \rightarrow \infty$ and thus

$$(24) \quad \lim_{k^* \rightarrow \infty} V_B = \frac{-d}{1 - \delta}.$$

Finally, recall that ϵ^* is defined as the point where $G(\epsilon^*) = 0$ or equivalently where $V(R =$

$1|\epsilon^*) = V(R = 0|\epsilon^*)$. Since $\epsilon^* \geq \tilde{\epsilon}$, $P(\epsilon^*) \leq \epsilon^*$ and thus $R^*(P(\epsilon^*)) = 0$. Replacing $V^*(P(\epsilon^*))$ in $G(\epsilon^*)$ yields the following indifference condition:

$$(25) \quad d = (1 - \delta)[\epsilon^*(\delta V_G - \delta V_B + 1) - \delta V_G] + \delta(\delta V_G - \delta V_B + 1)[\epsilon^* - P(\epsilon^*)].$$

Since $\epsilon^* - P(\epsilon^*)$ converges to zero as $\epsilon^* \rightarrow \tilde{\epsilon}$, regulation is fully persistent if:

$$(26) \quad \frac{d}{1 - \delta} \leq [\tilde{\epsilon} + (1 - \tilde{\epsilon})\kappa] \left[1 + \delta \frac{d}{1 - \delta} \right]$$

The left hand side of this equation represents the cost of permanent regulation. The right hand side represents the expected cost of deregulating in the steady state and then permanently regulating once the bad state occurs. Solving for d and bringing this result together with the foregoing discussion leads to the strategy outlined in the proposition.

When $d > \rho_{BB}$ regulation is never static optimal, even after the most pessimistic beliefs that could arrive in equilibrium, so it will never be used. The argument for the $\tau\rho_{GB}$ cutoff for permanent regulation is given in the text. We know beliefs fall over time from $\epsilon = 1$ to $\tilde{\epsilon}$ via the Markov process, so the characterization in Lemma 1 gives the result for the intermediary case. ■

PROPOSITION 4:

Proof. All the results follow from Proposition 2 except the derivation of the new cut-off τ' . Clearly, if experimentation is not used in equilibrium, the planner's optimal strategy will look identical to that in Proposition 3, so we will limit our attention to the cases where experimentation is used prior to deregulation. At the cutoff ϵ^{**} , the planner is indifferent between experimentation and regulation, so

$$(27) \quad \delta V^{**}(P(\epsilon^{**})) = -c + [\epsilon^{**} + (1 - \alpha)(1 - \epsilon^{**})][\delta V^{**}(P(\hat{\epsilon}^{**}))] + (1 - \epsilon^{**})\alpha[\delta V^{**}(P(0))].$$

Just as in the model without experimentation, $V^{**}(P(0)) = \kappa[-1/\delta + V^{**}(P(1))]$, and as ϵ^{**} approaches $\tilde{\epsilon}$ from above, $k^{**} \rightarrow \infty$ and thus

$$(28) \quad \lim_{k^{**} \rightarrow \infty} V_B = -d/(1 - \delta).$$

At this limit, equation (27) becomes

$$V = -c - \kappa(1 - \tilde{\epsilon})\alpha + [\tilde{\epsilon} + (1 - \tilde{\epsilon})(1 - \alpha)]V + \kappa(1 - \tilde{\epsilon})V,$$

where $V = \frac{-\delta d}{1 - \delta}$. Replacing for V and solving for d yields the cutoff in Proposition 4. ■