2014

On Experimentation and Real Options in Financial Regulation

Matthew L. Spitzer
Northwestern University Pritzker School of Law, matthew.spitzer@law.northwestern.edu

Eric L. Talley
Columbia Law School, etalley@law.columbia.edu

On Experimentation and Real Options in Financial Regulation

Matthew Spitzer and Eric Talley

ABSTRACT

Financial regulators have recently faced enhanced judicial scrutiny of their cost-benefit analysis (CBA) in advance of significant reforms. One facet of this scrutiny is judicial skepticism toward experimentation (and the real option to abandon) in the CBA calculus. That is, agencies have arguably been discouraged from counting as a benefit the value of information obtained through adopting new regulations on a provisional basis, with an option to revert to the status quo in the future. We study field experimentation versus more conventional forms of CBA (or analytic learning) in a regulatory-judicial hierarchical model. We demonstrate that there is no principled basis for dismissing (or demoting) experimentalism and that such rationales deserve a place in agencies' standard CBA arsenals. Nevertheless, our analysis also reveals an institutional reason for the tension between the judiciary and regulators, suggesting that regulators are plausibly too eager to embrace field experimentation while judges are simultaneously too recalcitrant.

1. INTRODUCTION

Thus far, the 21st century has proven decidedly unkind to the Securities and Exchange Commission (SEC). The commission has been targeted for considerable blame for its conduct leading up to the financial crisis,

MATTHEW SPITZER is the Howard and Elizabeth Chapman Professor of Law at Northwestern University. ERIC TALLEY is the Rosalinde and Arthur Gilbert Professor of Law, Business and the Economy at the University of California, Berkeley. Many thanks to Robert Bartlett, Andrew Daughety, Dan Farber, Jeff Gordon, Prasad Krishnamurthy, Alex Lee, Eric Posner, Joel Watson, Ivo Welch, conference participants at the 2013 Sloan conference Benefit-Cost Analysis of Financial Regulation at the University of Chicago Law School and the 2013 Theoretical Law and Economics conference at Vanderbilt Law School, and seminar participants at Iowa State University and the University of Chile for helpful comments. Samantha Strimling provided excellent research assistance. All errors are ours.

[*Journal of Legal Studies*, vol. 43 (June 2014)]

it has been embarrassed by its negligence in exposing infamous financial scandals (such as Bernard Madoff), and it has been overrun with the monumental tasks of implementing the two most comprehensive financial market legislative overhauls of the last 80 years (the Sarbanes-Oxley Act of 2002 and the Dodd-Frank Act of 2010). Arguably, however, none of these burdens has been as monumental as the epic "beat-down" the SEC has suffered in the courtroom, where it has encountered spirited (if not unprecedented) challenges to its rule-making authority under the Administrative Procedures Act (APA). In the last decade, the D.C. Circuit has applied the arbitrary-and-capricious standard under the APA to invalidate at least three significant SEC rule-making decisions (*Chamber of Commerce of the United States v. SEC* [412 F.3d 133 (D.C. Cir. 2005)]; *American Equity Investment Life Insurance Company v. SEC* [613 F.3d 166 (D.C. Cir. 2010)]; *Business Roundtable v. SEC* [647 F.3d 1144 (D.C. Cir. 2011)]).

A consensus poster child for the SEC's ministerial malaise is the 2011 case *Business Roundtable v. SEC*, where a three-judge panel for the D.C. Circuit invalidated a significant reform to the proxy rules that govern public companies. Under rule 14a-11, publicly traded issuers and investment companies would have been required (under certain circumstances) to allow shareholders controlling at least 3 percent of an issuer's voting securities to place their own nominees on the ballot for regular directorial elections. (This would have represented a mandatory change from the status quo ante, in which dissident shareholders must usually underwrite a proxy challenge, often at considerable risk and expense.)

After a contentious, party-line vote by the commission approving the rule change, the Business Roundtable and U.S. Chamber of Commerce filed a timely (and nearly immediate) challenge to the reform. Writing on behalf of a unanimous three-judge panel, Judge Douglas Ginsburg held that the commission had acted arbitrarily and capriciously in failing to assess the economic effects of the new rule, and accordingly the court invalidated and vacated the rule change. In particular, the court held that the SEC inconsistently and opportunistically framed the costs and benefits of the rule, that it failed to respond to potential problems raised in the notice and comment period, that it failed to quantify adequately certain costs or otherwise explain why they could not be quantified, and that it neglected to support its predictive judgments about the rule.[1]

1. The plaintiffs did not challenge—and the court did not overrule—a contemporaneous rule change under rule 14a-8, which prevents companies from excluding shareholders' proposals to establish procedures for proxy access.

We are personally uncertain whether—had it survived judicial scrutiny—rule 14a-11 would have generated positive or negative net economic effects. It is unclear, in fact, whether *anyone* could—in good faith—conjure a convincing case for either position. Indeed, an issue that continually plagues empirical corporate governance research is the challenge of using observational studies to demonstrate much of anything, much less the likely effects of novel reforms.

Rather, our interest in *Business Roundtable* (and cases of its ilk) focuses precisely on its epistemic indeterminacy: how, if at all, should one think of cost-benefit analysis (CBA) in situations in which theory is contested, empirics unclear, and politics pervasive? In this paper, we advance the argument that many of the recent judicial opinions related to financial regulation have placed too much emphasis on the ex ante empirical quantifiability of the costs and benefits of a proposed rule, giving short shrift to the role of regulatory experimentation (with an embedded real option to abandon) as a part of a CBA. Specifically, we argue that the judiciary's rhetoric in these recent cases effectively deters regulators from touting as one of the benefits from regulation the information produced through field-testing a new rule and the concomitant option to revert to the status quo ante (if appropriate) should the test prove unsuccessful. The failure to appreciate the real-option value of regulatory experimentation, we argue, has induced courts to be overly skeptical of innovative administrative reforms. Moreover, as financial markets grow more complex, nuanced, and difficult to measure with observational data, the real-option value of regulatory field experimentation also becomes greater.

To investigate and illustrate our argument, we develop and analyze a game-theoretic framework of an administrative-judicial hierarchy, using it to assess the relative importance of traditional CBA (which we call "analytic learning") versus the learning-by-doing benefits of regulatory experimentation. Using the model, we demonstrate that indeed there is no principled a priori basis for categorically disfavoring (or favoring) field experimentation within a cost-benefit framework. That is, depending on the facts and circumstances of a given case, and the technologies for regulatory learning available, an optimal regulatory policy might rely exclusively on analytic learning, exclusively on field experimentation, or on some combination of the two.

Nevertheless, our analysis also suggests that there may be a structural reason for the evident tension between the SEC's experimentalist zeal and the D.C. Circuit's squinty-eyed skepticism—one that does not necessarily hinge on differing ideological commitments among regulators and judges. Rather, our model suggests that the relative costs and benefits of each type of regulatory learning are not evenly distributed among the relevant players. Regulators—who disproportionately bear the costs of analytic learning—are likely too eager to embrace field experimentation, while the judiciary—who disproportionately bear the costs of experimentation—are simultaneously too recalcitrant. In the end, we argue, neither unfettered regulatory license nor unfettered judicial veto rights are likely to give rise to an optimal scheme for regulatory learning.

Three caveats deserve mention before proceeding. First, although we limit our arguments in this paper to financial regulation (focusing on the corporate and securities context), our analysis would potentially be applicable to other forms of regulation, such as environmental, communications, transportation, tax, and the like (see, for example, Sabel and Simon 2011; Sunstein 2002). In each case, the relative merits of experimental versus analytic learning would hinge on the relative costs, benefits, and precisions of each form of learning. Robust capital markets stand out, however, because participants regularly price out regulatory reforms in observable ways, thereby providing rapid and probative feedback to regulatory decision makers. By comparison, other domains—such as environmental regulation—plausibly require more protracted experiments yielding more recondite results. Given the amenability of securities markets to field testing, then, it is perhaps ironic that courts have arguably scrutinized the SEC's experimental efforts with greater (not lesser) skepticism than other regulatory domains.

Second, this paper is far from the first to suggest the use of experimentation in financial regulation. Even before the Business Roundtable complaint was filed, one of us advanced the thesis that rule 14a-11 was best viewed as a field experiment, yielding valuable future information about whether rule change would be worth keeping (Talley 2010). This debate joins a larger one about whether administrative agencies have anything near the requisite expertise and knowledge to regulate effectively and efficiently, particularly when randomized, clinical experimentation is impractical (see, for example, Romano 2012; Whitehead 2012; Posner and Weil 2013; Forstall 2012; McGinnis 2013).

Perhaps closest to our analysis in this regard are recent papers by Lee (2013) and Gubler (2014), both of whom also criticize the *Business Roundtable* decision while advocating greater judicial deference to the use of field experiments. Our approach joins with that of Lee and Gubler in this assertion. Our approach goes further, however, in that we situate our analysis in a game-theoretic institutional setting involving regulators and judicial actors who have not only policy preferences, but also specific preferences about how costs are distributed in a regulatory hierarchy. In so doing, we show that while field experimentation continues to have clear value in this setting, the players' places in the regulatory process may cause regulatory experimentation to be overvalued by the regulator (even as it is excessively discounted by the judiciary).[2]

Finally, in framing our analysis, we use *Business Roundtable* as a salient example of judicial hostility toward administrative experimentation.[3] At the same time, it bears noting that although numerous commentators at the time debated rule 14a-11 in terms of its experimental value (Talley 2010; Ribstein 2010), the SEC's own CBA tended to downplay that perspective. It is an open question whether—had the CBA more forcefully embraced experimentalism—the rule would have survived judicial scrutiny. On the basis of our reading of the case (as well as other anecdotal evidence),[4] we have reason to be skeptical, and we proceed on that basis.

In any event, other recent judicial opinions have suggested that the crescendo of judicial scrutiny over financial regulation has perhaps begun to abate. In the recent case of *Investment Co. Institute v. Commodity Futures Trading Comm'n* (720 F.3d 370 [D.C. Cir. 2013]), for example, the D.C. Circuit upheld the Commodity Futures Trading Commission's rule change requiring public registration by investment companies that were previously exempt. Many of the arguments advanced by the plaintiffs (indeed many of the plaintiffs themselves) were the same as in *Business Roundtable*. Although one case does not make for a trend, it at the

2. Sunstein (2014) recommends the use of break-even analysis—subject to a maximin criterion—to inform cost-benefit analysis (CBA) in financial regulation in information-impoverished environments. Although his analysis is largely silent on the question of regulatory experimentation and the real option to abandon, the framework we propose nests comfortably within Sunstein's proposal.

3. Others have similarly observed that the case marks an increasing escalation in the Federal Circuit of scrutiny over financial regulation. (For example, we note that many other commentators have made similar observations about the case; for example, Cox and Baucom [2012]; and Kraus and Raso [2013].)

4. Shortly after the publication of *Business Roundtable*, at an endowed lecture delivered by Judge Ginsburg related to the case, coauthor Talley (acting as commentator) posed the question of whether rule 14a-11 might have been more fruitfully assessed through the lens of field experimentation. Judge Ginsburg's response suggested significant resistance to regulatory experimentalism, going so far as to draw an analogy to the infamous U.S. Public Health Service Tuskegee syphilis experiment (1932–72). Viewed in this light, the lack of explicit reliance on experimentalism in the CBA for rule 14a-11 was plausibly the byproduct of (correctly) anticipated hostility toward such rationales.

very least suggests that the D.C. Circuit has begun to reconsider the appropriate role of regulatory experimentation. Nevertheless, with the judicial situation in flux, and with Congress currently entertaining legislation that would require enhanced CBA for financial regulation,[5] this is a particularly apt time for academic participants to make meaningful contributions to the debate.

Our analysis proceeds as follows. Section 2 develops a dynamic game-theoretic framework for our analysis, including the formal development of two different (and nonexclusive) modalities for regulatory learning. Section 3 analyzes how the judiciary and administrative agencies interact in a judicial and/or regulatory hierarchy. Here we demonstrate that various institutional factors can induce regulators to be (inefficiently) too anxious to experiment while making judges (inefficiently) too recalcitrant. The section also demonstrates that the interaction of regulators and the judiciary during the judicial review process does not necessarily improve things from a welfare perspective. Section 4 discusses the robustness of our results along with several possible extensions. Section 5 concludes.

2. FRAMEWORK AND PRELIMINARIES

In this section we develop and analyze a game-theoretic framework for assessing the relative benefits of conventional empirical cost-benefit research versus the benefits of field experimentation. Although based loosely on *Business Roundtable v. SEC* (647 F.3d 1144), the framework described is deliberately abstract, and—as noted above—it may be generalizable to other areas involving regulatory learning in administrative law. The information structure is similar in spirit to the framework developed in Spitzer and Talley (2000, 2013), but it is adapted for the specific context of the problem at issue here.

Consider an economy of size N with a time horizon of $T + 1$ discrete periods, indexed by $t \in 0, 1, 2, \ldots, T$. Period $t = 0$ denotes an ex ante period. For each period $t \geq 1$, the economy is regulated by a single policy within a policy space $Y \equiv \mathbb{R}^1$, where element $y \in Y$ represents a specific policy. The implied ordering of the policy space is deliberate, and it lends itself to many real-world policy debates; for the sake of concreteness, suppose that the policy decision reflects the impediments placed on minority shareholders of public companies who wish to nom-

5. See, for example, Financial Regulatory Responsibility Act of 2013, S. 450, 113th Cong. (2013). For a good overview, see Bartlett (2014).

inate their own candidates in corporate board elections. Under this interpretation, larger values of y correspond to increasingly more conservative policy stances (that is, unfriendly to minority shareholders). The reverse movement corresponds to increasingly liberal stances (that is, solicitous of minority shareholders). Without loss of generality, we normalize the policy $y = 0$ to represent the maximally centrist position in policy space.

Suppose that at $t = 0$, a default status quo policy is already in place and that its location is common knowledge to all participants (for example, by dint of extensive experience with the incumbent policy). Denote y_0 as the location of the status quo, and further assume (also without loss of generality) that this location takes on a positive value,[6] so $y_0 = \alpha > 0$.

Also at $t = 0$, suppose that an alternative policy (denoted y_1) has been proposed that would move the regulatory regime away from the status quo.[7] Unlike the status quo, however, we assume the alternative policy is unfamiliar and untested—a key challenge that typifies much of rule making in financial regulation. To capture this factor analytically, we assume that y_1's true location in policy space Y is not known with certainty. Rather, its location is commonly known to be a random variable (Y_1) distributed normally with mean μ and precision τ.[8] The uncertainty about the location of the alternative drives our analysis in two ways. First, the players are assumed to be averse to risk, and second, two technologies are available to the regulator that can generate more precise information about the distribution of Y_1. Both of these dimensions are detailed below.

Implementing a new policy imposes a lag of 1 period. Thus, if a switch to the alternative policy (or a switch back from the alternative) is declared at time t, the new policy will become effective in period $t + 1$.

There are three relevant players in the game: the public (P) of size N, which is ultimately regulated by either the status quo ante policy or

6. All the intuitions below follow (but from the other direction) when y_0 takes on negative values. The only thing necessary to impel our analysis forward is for the status quo ante to be distinct from the first-best choice of at least one of the relevant decision-making constituencies.

7. We do not consider (at least for now) how this alternative comes to the fore. However, the equilibrium behavior described below would characterize the continuation payoffs of an initial stage in which the policy alternative is selected.

8. Recall that the precision of a normal random variable is the reciprocal of the variance, so $Y_1 \sim N(\mu, [1/\tau])$. In the analysis that follows, it is more economical to express variance measures in terms of precision.

the alternative policy; a unitary regulator (R), who makes an initial decision about whether to retain the status quo or embrace the alternative (and on what terms); and a unitary judiciary (J), which is in a position to veto the regulator's choice to implement the alternative. The judiciary's veto right is asymmetric and takes hold only on the regulator's adoption of the alternative policy; that is, J cannot veto a regulator's choice to maintain the status quo.

2.1. Preferences and Welfare

Each player in the game receives payoffs that include (but need not be limited to) her preferences over regulatory policy.[9] Specifically, suppose that each player $i \in \{P, R, J\}$ has an ideal point y_i^* in policy space Y. These ideal points correspond to the policy locations that—if known with certainty—would be most suited to each player's ideological dispositions. For the public, this ideal point might correspond to the policy preference of a representative citizen or alternatively that of the median citizen or voter. For players R and J, this ideal point may coincide with the public's or may be determined in other ways. Denote these ideal points by $\{y_P^*, y_R^*, y_J^*\}$ for players {P, R, J}, respectively. We normalize the public's policy preference at $y_P = 0$. Thus, as with the policy space, if another player has a policy preference $y_i^* > 0$, we describe that player as being more conservative than the public; conversely, if player i has policy preference $y_i^* < 0$, she is more liberal than the public. We concentrate most of our analysis on the special case where, like the public, the regulator and judiciary share centrist policy commitments, so $y_R^* = y_J^* = y_P^* = 0$. This restriction not only simplifies our analysis, but it allows us to concentrate more squarely on the structure of the judicial-regulatory hierarchy. We discuss the implications of relaxing this assumption in Section 4.[10]

Each player derives utility in each period (at least in part) from minimizing the expected squared distance between its ideal point y_i^* and the governing policy during each period. (All players discount future periods' payoffs using a common discount factor $\beta \leq 1$.) The concave nature of each party's per-period utility function in policy space suggests that, ceteris paribus, the players will exhibit risk aversion in their policy pref-

9. In addition, the regulator also has preferences about uncompensated costs that it incurs in conducting certain types of CBA.

10. We are constrained by the length limitations for this symposium submission. We hope to extend the build-out model further in this direction in later papers.

erences in each period, inducing a bias toward the known quantity represented by the status quo ante.

Finally, in order to facilitate comparisons across different regimes, the model must commit to a welfare measure. In what follows, we equate social welfare to the weighted sum of the expected payoffs of the regulator, the courts, and the public, scaling the latter's payoff by the size of the economy, N. Consequently, as the economy grows, the relative importance of the regulator's and the judiciary's payoff shrinks proportionally (even though their incentives remain relevant to the ultimate choice of policy).[11]

2.2. Regulatory Learning

Because the true location of the alternative policy is stochastic, the risk-averse players value acquiring additional knowledge about the realized location of y_1. The administrative agency plays the key role in producing such information, and, accordingly, we consider two distinct (but not mutually exclusive) forms of regulatory knowledge acquisition: analytic learning and field experimentation. We describe each approach below.

2.2.1. Analytic Learning. This form of regulatory learning is perhaps the most familiar manifestation of what many traditionally envision when considering CBA by administrative agencies. Under analytic learning, the regulator attempts (at some cost) to harvest a variety of evidence (theoretical, empirical, anecdotal, experimental, and the like) that is probative of the alternative policy's characteristics. By so doing, the regulator (and others) gains more precise information about the alternative's desirability relative to the status quo. All else constant, in the light of the players' risk aversion, greater precision is valuable.

To capture the intuition of analytic learning, we assume that R can extract a signal (denoted Z) about the location of y_1 within the policy space. Specifically, signal Z is assumed to be normally distributed with a mean equal to the true value of y_1 and a precision of γ. We assume that R may choose any precision $\gamma \geq 0$ it desires. However, the regulator's choice of precision also imposes a direct cost of effort, $c(\gamma)$, which we

11. Because N is essentially a free parameter of the model, it may also be interpreted as reflecting broader considerations about how to trade off the public's welfare against governmental actors' welfare.

assume is uncompensated and borne solely by the regulator.[12] Specifically, suppose that $c(0) = 0$, $c'(\gamma) > 0$, and $c''(\gamma) > 0$, so increasing precision comes at an increasing marginal cost. (One possible example is a quadratic cost, $c(\gamma) = c_0(\gamma^2/2)$ for $c_0 > 0$, so the costs of precision increase in the square of the target precision level.) The fact that the regulator must bear the uncompensated costs of analytic learning can introduce important agency costs into the model.[13]

If the regulator engages solely in analytic learning, it must opt for either the status quo ante or the alternative immediately after observing Z, and it may not alter course thereafter up to period T.

2.2.2. Field Experimentation. A second form of learning available to the regulator is what we call field experimentation. Under this approach, the regulator effectively test drives the alternative policy, embracing it on a trial basis that lasts for K periods (where $K \leq T$). At the end of K periods, the regulator must choose whether to keep the alternative in place or to revert back to the original status quo ante (starting in period $K + 1$). Field experimentation is roughly tantamount to a policy change that includes a mandatory sunset provision so that the policy change must be reactivated once the trial period ends (see, for example, Gersen 2007).

To capture the intuition of field experimentation, we assume that in each period the alternative policy is in place, it generates probative data about the location of y_1. Specifically, for each period of the field test $k = 1, 2, \ldots, K$, a signal V_k about the alternative is generated. The sequence of signals is denoted $\{V_1, V_2, \ldots, V_K\}$, and we assume that each V_i is independently and normally distributed with mean equal to the true value of y_1 and precision ω. Once the experimental period is complete (that is, at the completion of period K), R must immediately choose between making the experiment permanent or returning to the status quo, and this decision becomes effective in period $K + 1$ lasting

12. The assumption that the costs of analytic learning are borne solely by the regulator can be relaxed, but many of our results hinge on the relative costs of analytic learning versus experimentation being larger for the regulator than for the judiciary. We justify this assumption with a more extended discussion in Section 4.

13. It is important to acknowledge that because regulators are nominally provided with a government budget, at least some portion of verifiable costs associated with analytic learning are compensated. That said, it is eminently plausible that regulators bear at least some uncompensated fraction of analytic learning costs, either because they are undercompensated for the work they do, or because budgets are arguably exogenous to workload, or because regulators must allocate budgets across multiple projects, generating a noteworthy opportunity cost.

until terminal time T. Unlike analytic learning, field experimentation does not impose an idiosyncratic cost on the regulator; however, it may still be perceived as costly to all parties to the extent that the status quo appears preferable given the information available at the time the experiment commences.

2.2.3. Hybrid Learning. The two learning approaches can be employed in isolation, but they can also be combined. That is, the regulator could choose to invest in some analytic learning ($\gamma > 0$) while also committing to engage in field experimentation ($K > 0$). Consequently, the regulator's learning plan entails two choice parameters: (γ, K). In order to keep things tractable (at least at this stage), we make a simplifying assumption that the real option to abandon the experiment is a "European" option, exercisable only at the conclusion of its maturity period K. In other words, the regulator must commit to regulatory learning policy (γ, K) at the onset at $t = 0$, and it is not free to alter its plan in response to interim signals that the experiment might yield while underway. (Equivalently, R may be unable to observe such signals until the trial period is complete.) Only when the experimental period ends may the regulator marshal all available evidence to make a decision about whether to continue embracing the alternative or revert to the status quo—a decision that endures thereafter until terminal time T. This constraint permits one to isolate the option value associated with regulatory learning at a single point, K.[14] Hybrid learning protocols create something akin to the learning that occurs in well-known bandit problems in decision theory (for example, Berry and Fristedt 1985), albeit one that is in this case nested within a strategic environment.

As a final aside, it bears noting that this specification is rich enough to allow for the consideration of most of the serious policy contenders that are related to regulatory learning. Specifically, the regulator can choose to engage exclusively in analytic learning (with no field experimentation) by fixing $\gamma > 0$ and $K = 0$. In contrast, the regulator can engage exclusively in field experimentation (with no analytic learning) by fixing $\gamma = 0$ and $0 < K < T$. Finally, the regulator can effectively short-circuit all learning—issuing an immediate, uninformed choice between the status quo and the alternative—by choosing $\gamma = K = 0$.

14. While this is a restrictive assumption, note that relaxing it would permit the regulator to exercise a real option at other junctures in the model, thereby enhancing the relative attractiveness of experimentation. Thus, the analysis below could be seen as identifying a lower bound for the value of experimentation coupled with the real option of abandoning.

2.3. Regulatory Hierarchy and Sequence of Moves

Finally, as noted above, this paper situates a comparison of experimental and analytic learning in a hierarchical model of regulatory-judicial interaction. We assume the following sequential interaction between R and J:

1. In period $t = 0$, R announces a learning plan (γ, K) associated with its CBA.
2. Immediately thereafter (also in period $t = 0$), R's plan is challenged by an interest group and brought before J. J makes a decision (upholding or overturning) that maximizes its expected payoff given its information and conjectures about equilibrium play at that stage of the game.
3. If J overturns the plan at $t = 0$, the status quo rule ante remains in effect through period T.
4. If J upholds, then R's plan goes into effect and y_1 is put into effect through period $t = K$, at which point R makes a posterior decision about whether to retain y_1 for the remaining periods or instead to revert to y_0.
5. Regardless of whether R decides at $t = K$ to retain or revert, its decision is once again immediately subject to review, and J may veto or uphold the agency's posterior decision.

Although this structure provides a relatively intuitive benchmark, it goes without saying that no extensive form is sacrosanct and that other sequential structures are plausible as well. (Our discussion takes on other candidates in Section 4, offering conjectures about their effects on our results.)

3. ANALYSIS OF MODEL

With this analytic framework in hand, we now proceed to analyze and assess the plausible equilibria that emerge. We proceed first by characterizing the benchmark socially optimal CBA, contrasting it to R's and the J's respective preferred plans (were they to have absolute decision-making authority). We then turn to analyzing how the judiciary and agency behave in equilibrium, comparing that outcome with the social optimum.

3.1. Socially Optimal Cost-Benefit Analysis

Consider a first-best learning protocol, in which (γ, K) are chosen by a benevolent social planner in order to maximize expected social welfare.

Recall that under the status quo, the policy location of y_0 is known with certainty ($y_0 = \alpha \geq 0$), and thus each player's expected payoff for each period that is associated with the status quo is the expected squared distance between the ideal point and α:

$$\mathrm{EU}_i(y_0) = E[-(y_i^* - y_0)^2] = -\alpha^2. \tag{1}$$

As of $t = 0$, imposing the assumption that all players share the ideal point of 0 in policy space, staying with the status quo yields an expected payoff for each player:

$$\mathrm{EU}_i(y_0) = \sum_{t=1}^{T} -\beta^t(0 - \alpha)^2 = -\alpha^2\left(\beta\frac{1 - \beta^T}{1 - \beta}\right). \tag{2}$$

Aggregating across time and all three constituencies, and recalling that the size and welfare weight of the public is N, total joint welfare associated with the status quo ante is given by

$$\mathrm{ESW}(y_0) = -\alpha^2\left(\beta\frac{1 - \beta^T}{1 - \beta}\right)(N + 2). \tag{3}$$

Now consider, in contrast, the payoffs associated with an announced regulatory learning plan (γ, K). Note that because all random variables and signals are assumed to come from the conjugate family of normal distributions, any distribution that is conditioned on combinations of random variables and observed signals is also normal. Consequently, suppose that the trial period has been completed, revealing both an analytic-learning signal Z and a series of field experimentation signals $\{V_1, \ldots, V_K\}$. The conditional random variable $(Y|Z, V_1, \ldots, V_K)$ is distributed normally with the following parameters:

$$(Y_1|Z,\ V_1,\ \ldots,\ V_K) \sim N\left(\frac{\tau\mu + \gamma z + \omega\Sigma_{t=1}^{K} v_t}{\tau + \gamma + \omega K},\ \frac{1}{\tau + \gamma + \omega K}\right). \tag{4}$$

To economize on notation, it is possible to compress R's observed information into the observation of a single, hybridized signal X,with realization

$$x = \frac{\tau\mu + \gamma z + \omega\Sigma_{t=1}^{K} v_t}{\tau + \gamma + \omega K}. \tag{5}$$

It is straightforward to confirm that the unconditional distribution of X is normal with mean μ and precision $\tau[1 + (\tau/\gamma + \omega \mathrm{K})]$. (See the Appendix for this and all other derivations and proofs.) Conditional on observing hybridized signal x, the expected payoff of choosing the alternative policy is

$$-E([Y_1 - y_i^*]^2|x) = -E([Y_1 - 0]^2|x) = -(E[Y_1|x])^2 - \mathrm{Var}(Y_1|X)$$
$$= -(x)^2 - \frac{1}{\tau + \gamma + \omega K}.$$

Consequently, all parties will favor retaining the alternative strategy if and only if its per-period expected payoff from the alternative strategy (conditional on CBA learning) exceeds its known payoff under the status quo ante. In other words, the planner will favor the alternative after observing x if and only if

$$\alpha^2 > x^2 + \frac{1}{\tau + \gamma + \omega K}. \quad (6)$$

In principle, the hybrid signal may be so far from the social optimum that this condition is not satisfied. In fact, the condition may not be satisfied even when the hybrid signal conveys good news about the estimated location of the alternative (that is, that $x = 0$). This latter intuition is formalized in the following lemma.

Lemma 1. It is never socially optimal to remain with the alternative strategy after experimentation if $\tau + \gamma + \omega K < 1/\alpha^2$.

The condition in lemma 1 is a sufficiency condition for the status quo to remain socially optimal regardless of what information is generated from the CBA learning plan (γ, K). Effectively, it states that no amount of regulatory effort can render the alternative an attractive option if the aggregate information derived from that effort is too imprecise. If the condition in lemma 1 is not satisfied, then it may be optimal to embrace the alternative policy over the status quo after experimentation, but only if the hybrid signal X falls into the interval

$$x \in [-x^*(\gamma, K), x^*(\gamma, K)], \quad (7)$$

where $x^*(\gamma, K) = +\sqrt{\alpha^2 - (1/\tau + \gamma + \omega K)}$. In other words, the hybrid signal X must be within a symmetric interval around the players' common ideal point to justify embracing the status quo and forgoing experimentation.

Combining these considerations, the social welfare problem is to choose a learning strategy (γ, K) that maximizes the expected improvement in social welfare over the status quo, or $\Delta_{\mathrm{ESW}}(\gamma, K) \equiv \mathrm{ESW}(\gamma, K) - \mathrm{ESW}(y_0)$. In terms of the model's parameters, this measure is defined as follows:

$$\Delta_{\text{ESW}}(\gamma, K) = -c(\gamma) - (N+2)\left(\beta\frac{1-\beta^{K}}{1-\beta}\right)\left(\mu^2 + \frac{1}{\tau} - \alpha^2\right) \\ - (N+2)\beta^{K+1}\frac{1-\beta^{(T-K)}}{1-\beta}E_x\left[\min\left\{0,\ x^2 + \frac{1}{\tau+\gamma+\omega K} - \alpha^2\right\}\right], \tag{8}$$

where the first term on the right-hand side is the social cost of analytic learning, the second term is present discounted value (PDV) of the social cost of experimentation, and the third term is PDV of the social value of the real option to abandon.

More specifically, the above total social welfare measure has three component parts. The first component is the regulator's direct cost of engaging in the analytic learning, or $c(\gamma)$. The second is the aggregate present value of the expected social welfare loss associated with the experimental period, where the status quo ante (of $-\alpha^2$) is passed up for K periods in exchange for the expected (risk-adjusted) payoff of the alternative (or $-[\mu^2 + (1/\tau)]$). (Given the assumptions in the lemma, this cost to social welfare of experimentation is always positive.) The third component in expression (8) is the present value of the real option to abandon the reform at the end of the prescribed sunset period, informed by the fruits of the experimental plan. Note that this term enters $\Delta_{\text{ESW}}(\gamma, K)$ in a strictly positive fashion. Let $(\gamma^{\text{ESW}}, K^{\text{ESW}})$ denote the optimal learning plan in the sense of maximizing $\Delta_{\text{ESW}}(\gamma, K)$.

Analysis of the social planner's problem yields the following proposition (proof in the Appendix).

Proposition 1. If $\alpha^2 > 1/\tau$ and $c'(0)$ is sufficiently small, then $\gamma^{\text{ESW}} > 0$. Similarly, if $\alpha^2 > 1/\tau$ and μ is sufficiently small, then $K^{\text{ESW}} > 0$.

The intuition behind proposition 1 is simple and intuitive. It essentially states two results: First, it states that some analytic learning will be socially optimal so long as the marginal cost of investing in precision (captured by γ) is sufficiently low and the location of the status quo ante is not too close to first best. Second, it states that at least some experimentation will be socially optimal so long as the cost of experimentation (captured by μ) is sufficiently low and the status quo ante is not too close to first best.

Note that the sufficiency conditions given in proposition 1 are similar but not identical for the inclusion of analytic learning and field experimentation (respectively) in an optimal CBA. In other words, an implication of proposition 1 is that the social optimum might involve a corner solution consisting solely of either analytic learning with no experimentation or experimentation with no analytic learning. However, there does not appear to be any categorical reason to expect one of these corner solutions to obtain rather than the other. Moreover, the conditions in proposition 1 may be satisfied simultaneously for both analytic learning and field experimentation. This result is sensible, even if it is inconsistent with the view (arguably harbored by at least some recent judicial opinions and proposed legislation) that field experimentation should be categorically disfavored against other forms of CBA. To the contrary, proposition 1 suggests that the importance of field experimentation relative to analytic learning turns on the facts and circumstances of the case.

3.2. Institutional Preferences

Before proceeding to the strategic interaction between the two strategic players R and J, we first consider their preferences over policy generally—that is, if each had the absolute right to implement his or her most preferred policy. As before, we continue to assume that the public, the regulator, and the judiciary share the same ideal point in policy space; nevertheless, the institutional structure of their interaction suggests that an agency cost problem can exist because of the distribution of social costs and benefits associated with CBA.

Consider first the net payoff to the agency associated with a learning plan (γ, K), which we denote $\Delta_R(\gamma, K)$:

$$\begin{aligned}\Delta_R(\gamma, K) = & -c(\gamma) - \left(\beta\frac{1-\beta^K}{1-\beta}\right)\left(\mu^2 + \frac{1}{\tau} - \alpha^2\right) \\ & - \beta^{K+1}\frac{1-\beta^{(T-K)}}{1-\beta}E_x\left[\min\left\{0,\ x^2 + \frac{1}{\tau+\gamma+\omega K} - \alpha^2\right\}\right].\end{aligned} \quad (9)$$

Although this expression looks similar to the social welfare measure above, it differs in a few important ways: in equation (9) the regulator internalizes neither the full social cost of experimentation (that is, that portion borne by the public or the judiciary) nor the full social benefit of the real option to abandon. However, the regulator does internalize the full social cost of investing in analytic learning (that is, $c(\gamma)$). Consequently, if the regulator were left to its own devices to maximize $\Delta_R(\gamma, K)$, it would tend to oversupply experimentation and undersupply analytic learning in its CBA.

Stated more formally, comparing equation (9) with equation (8) yields the following propositions.

Proposition 2. The regulator (weakly) prefers to engage in too little analytic learning and excessive field experimentation relative to the social optimum. However, the regulator would implement a socially optimal learning protocol if it were constrained to choose any $\gamma \geq \gamma^{\mathrm{ESW}}$.

One notable feature from proposition 2 is the observation that the regulator can be induced to adopt a socially optimal research plan solely by requiring it to invest in at least as much analytic learning as is characterized by γ^{ESW}. In other words, once the regulator is constrained to engage in no less than the socially optimal amount of analytic learning, it would proceed to select the first-best level of experimentation on its own accord. This lower bound might be interpreted as the analog of a first-best arbitrary-and-capricious standard.

Unfortunately, it is unlikely that the judiciary would adopt this standard on its own. Indeed, consider the expected net payoff for the judiciary associated with a learning plan (γ, K), which we denote $\Delta_{\mathrm{J}}(\gamma, K)$:

$$\begin{aligned}\Delta_{\mathrm{J}}(\gamma, K) = &-\left(\beta\frac{1-\beta^{K}}{1-\beta}\right)\left(\mu^2 + \frac{1}{\tau} - \alpha^2\right)\\ &- \beta^{K+1}\frac{1-\beta^{(T-K)}}{1-\beta}E_x\left[\min\left\{0,\ x^2 + \frac{1}{\tau+\gamma+\omega K} - \alpha^2\right\}\right].\end{aligned} \tag{10}$$

As with the regulator, the judiciary fails to internalize a portion of the experimental social cost as well as a portion of the real option to abandon. However, because neither the judiciary nor the public bears the direct investment costs associated with analytical learning, its maximand is identical (up to a scalar) to that of the public. The judiciary's failure to internalize any of the costs of analytic learning will cause it to be too dismissive of experimentation and too anxious to promote pure analytic learning. This reasoning is captured by the following proposition:

Proposition 3. The judiciary prefers to engage in (weakly) too much analytic learning and (weakly) too little field experimentation relative to the social optimum.

Just as proposition 2 provides reason to be skeptical of granting regulatory actors unfettered license to balance analytic learning against field experimentation, proposition 3 states that the judiciary suffers from the opposite problem: it has little interest in field experimentation, and it would rather attempt to induce the regulator to engage predominantly (or even solely) in analytic learning.

3.3. Equilibrium

Having considered both the socially optimal CBA and the institutionally induced preferences of the strategic players (R and J), we now proceed to consider how the parties would interact within a Bayesian-perfect equilibrium. We solve the game backwards, considering first the "sunset" stage in period $t = K$, assuming that a hybrid plan of the form (γ, K) has been installed by R and not overturned by J. Analysis of the players' incentives at this stage immediately yields the following lemma:

Lemma 2. Upon the termination of the experimental period for a learning plan (γ, K) in period K, the player R will opt to retain the alternative policy if and only if

$$x \in \left[-\sqrt{\alpha^2 - \frac{1}{\tau + \gamma + \omega K}},\ \sqrt{\alpha^2 - \frac{1}{\tau + \gamma + \omega K}}\right].$$

Player J will uphold whichever decision player R makes, and both players' decisions will be welfare maximizing as of period $t = K$.

The intuition behind lemma 2 is straightforward. At period K, the experimentation period has come to a close and all of the previous costs of analytic learning and experimentation are now sunk. Consequently, the only issue left to decide is which policy choice to make (given current information), and R will choose to retain the experimental alternative so long as the best posterior estimate of y_1's location (embodied by the realization of x) is sufficiently close to R's ideal point—an assessment corresponding to the condition stated in lemma 2. Moreover, by assumption, there is perfect alignment among R's, J's, and P's preferences in policy space, and thus J can do no better than to uphold R's decision at that stage. (Conflict between J and R could occur at this stage if they have different ideal points in policy space—a possibility we take up in Section 4.)

Moving backward in time to the ex ante period $t = 0$, both players will anticipate the behavior described in lemma 2 and will incorporate that into their strategy at the time the CBA is announced. Consider first the judicial actor J who is ruling on whether to invalidate R's learning plan (γ, K). Under the extensive structure described above, invalidation of the plan implies that the status quo ante will remain in place thereafter. Consequently, J could conceivably favor upholding R's announced learning plan even if that plan is not J's most preferred plan (that is, even if J would have preferred a weightier dose of analytic learning over field experimentation). Equivalently, the judiciary will veto the regulator's plan only if the CBA yields a payoff that is worse (from the judiciary's perspective) than simply maintaining the status quo for the duration of the game. Formally, this reasoning implies that J will uphold any learning plan (γ, K) such that $\Delta_J(\gamma, K) \geq 0$.

Anticipating J's behavior, R will design its learning plan to maximize its expected payoff, subject to the constraint that J must expect a nonnegative payoff going forward. That is, anticipating equilibrium play, R's choice at $t = 0$ boils down to selecting a learning plan (γ, K) to solve the following constrained optimization problem:

$$\max_{(\gamma, k)} \Delta_R(\gamma, k), \qquad \text{subject to} \quad \gamma \geq 0,\ K \geq 0, \text{ and } \Delta_J(\gamma, k) \geq 0. \tag{11}$$

As shown, if completely unconstrained, R would favor a learning plan that involves too much experimentation (and too little analytic learning) relative to the social optimum, while the judiciary would have an incentive to do the opposite. The constraint of eventual judicial review (reflected through the treatment of J's payoff as a constraint in the above optimization problem) raises the possibility that R's rule-making endeavors might be nudged toward the social optimum.

Somewhat surprisingly, however, it turns out that this plausible balance does not emerge from our framework, as reflected in proposition 4:

Proposition 4. There is a unique Bayesian perfect equilibrium in the sequential game defined above. In it, the regulator announces a learning plan (γ^R, K^R) that maximizes its own welfare $\Delta_R(\gamma, K)$ unconstrained by the possibility of judicial review. As per proposition 2, this plan entails (weakly) too little analytic learning and (weakly) too much field experimentation relative to the social optimum. This policy is never overturned by the judiciary.

Although the result in proposition 4 is somewhat surprising, its intuition is relatively straightforward. It turns out that the need to satisfy J's preferences is never a binding constraint in the modeling framework described above. To see why, suppose for argument's sake that the regulator simply ignored the judiciary's veto right and selected a learning plan (γ^R, K^R) so as to maximize its own personal payoff $\Delta_R(\gamma, K)$. Because this plan is optimal for the regulator, it must deliver R a nonnegative payoff over the status quo, and thus $\Delta_R(\gamma^R, K^R) \geq 0$. However, comparison of equation (9) with equation (10) makes it is clear that because only player R bears the total social cost of analytic learning, it must be the case that $\Delta_J(\gamma^R, K^R) \geq \Delta_R(\gamma^R, K^R)$. In other words, while J is not happy with R's favored mixture of analytic learning and experimentalism, R's choice is never so objectionable from J's perspective to justify maintaining the status quo. Consequently, then, at least within this institutional structure, the regulator acts just as in it did in proposition 2—as though it has an unfettered right to declare a policy without any judicial review.

4. ROBUSTNESS AND EXTENSIONS

The analyses from the previous sections help underscore three core insights of this paper. First, regulatory experimentation deserves to be considered alongside analytic learning as a bona fide means for conducting CBA—a conclusion that is arguably at odds (or at least in tension) with several recent judicial decisions and pending legislation. Second, notwithstanding the utility of experimentation in information-poor regulatory environments, institutional factors can cause regulators to be too zealous about experimentalism while causing judges to be too skeptical. Third, although the institutional interdependence of regulators and the judiciary through judicial review might—in principle—cause the players to moderate their preferences, such a result is not guaranteed, and indeed in our baseline model, judicial review provides a poor constraint on regulators' experimentalist zeal.

As with any model, of course, the analysis above has made several simplifying assumptions, the alteration or relaxation of which could conceivably bear on the above insights. We briefly explore a variety of these variations below.

First, the analysis has assumed that the costs of analytic learning fall disproportionately on the regulator rather than on the judiciary (or society). Although this assumption is not necessary to drive proposition 1 (regarding the social optimality of analytic learning and experimentation), propositions 2, 3, and 4 (regarding the difference in incentives between the agency and judiciary) all require this assumption. We believe that this assumption is justified on several grounds. Most directly, because analytic learning requires active research and synthesis by the agency, it is reasonable to believe that the marginal costs of such effort are not compensated by agency budgets.

More subtly, regulators may be more averse to the costs of analytic learning because they tend to serve for a much shorter expected tenure than do federal judges. Figure 1 illustrates the average tenures of sitting Federal Communications Commission and SEC commissioners relative

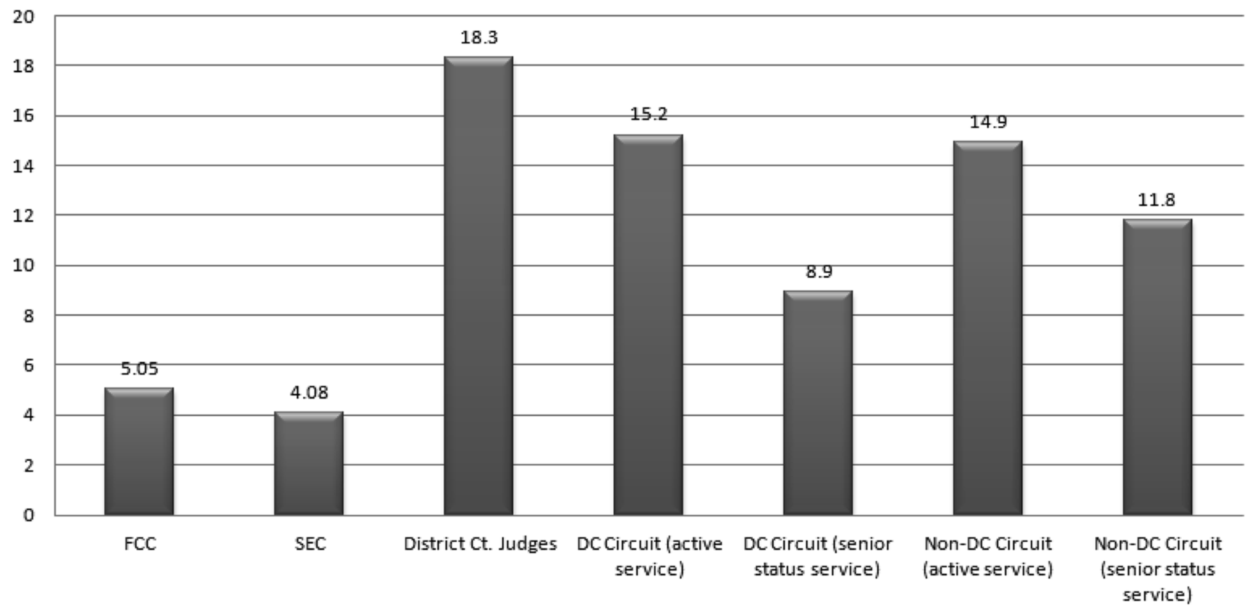

Figure 1. Regulator and federal judiciary tenures: average length of service (in years)

to federal district court judges, D.C. Circuit judges, and non–D.C. Circuit judges.[15]

The significantly shorter service of regulators implies that such actors are not able to amortize the costs of analytic learning over as long a period as the judiciary. Moreover, they generally need not confront the total cost of experimentation, since commissioners may well be on to another job by the time the experiment runs its course, while the federal judges remain sitting. These figures perhaps give firmer grounds for the model's assumption that regulators place greater relative weight than the judiciary on the costs of analytic learning over experimentation.

Second, our model has assumed that J was unable to commit to an adjudication strategy ex ante (such as requiring a minimal level of analytic learning by R); rather, J was constrained either to approve or to reject R's CBA in a manner consistent with J's payoff at the time of its decision. In equilibrium, J's inability to commit introduced a relatively lax constraint on R's behavior (see proposition 4), effectively permitting R to implement its preferred policy unconstrained by J. One possible variation of the model would be to assume that J has the means to precommit to rejecting any learning plan whose analytic learning component falls below some judicially determined threshold $\hat{\gamma}^{J}$. If this commitment level were credible, then R would be forced to reckon with a real judicial review constraint. Most optimistically, suppose J selected the socially optimal level of analytic learning (so $\hat{\gamma}^{J} = \gamma^{ESW}$); it is easy to show that R's optimal behavior in such a setting would be to announce a learning plan coinciding with the social optimum whenever it chose to move forward with regulation, so $(\gamma, K) = (\gamma^{ESW}, K^{ESW})$.[16] One po-

15. The data from Figure 1 come from the following sources. For Federal Communications Commission and Securities and Exchange Commission tenures, we consulted the agencies' websites to obtain lists of former members of each commission. Starting with appointments in 1961, we calculated the number of years in office for each commissioner. We did this for every commissioner to the present, omitting sitting commissioners, since we do not know how long they will serve. The data on district court judges come from Yoon (2003). For the data on circuit courts of appeal, see Federal Judicial Center, History of the Federal Judiciary, Export of All Data in the Biographical Directory of Federal Judges, 1789–present (http://www.fjc.gov/history/home.nsf/page/export.html), starting in 1960 and through the present for judges who are not still active. (For the non-D.C. circuits, we excluded the Federal Circuit because its docket does not include appeals from administrative agency rule making.)

16. We should note that it is possible under this set of assumptions that R would decide simply not to go forward with a rule change, even when it would be socially optimal to do so and when R's learning plan would be optimal. This is because R's policy payoff may not justify the uncompensated costs of analytic learning.

tentially charitable interpretation of cases such as *Business Roundtable* is that they represented an effort by the judiciary to alter equilibrium play prospectively, establishing precedents that would serve as a means for precommitment in later cases to a more demanding standard on analytic learning than regulators would otherwise adopt.

That said, there is no guarantee that even if J could precommit to a threshold value of γ, the judiciary would have the incentive to set it in a welfare-maximizing way. Rather, the judiciary might well simply maximize its expected payoff by fixing $\hat{\gamma}^J$ at its own preferred point γ^J, which (as shown above) entails more analytic learning than the social optimum. When J precommits to a threshold analytic-learning requirement in this way, the resulting equilibrium learning plan would now be optimal from J's perspective rather than from R's, but it need not be any closer to the social optimum.[17]

Another potentially interesting (and intuitively attractive) extension of the model would be to allow J to observe the fruits of R's analytic learning before conducting its review of R's plan. Under this alternative approach, J has an informational advantage when it acts—one that in some ways serves the interests of all players. Specifically, suppose J observed the analytic-learning signal Z before issuing its opinion, allowing it to bring more information to bear in its judicial review than R had in its original rule making. Although this variation adds some degree of technical complication to the model, a few important intuitions nonetheless emerge. First, the resulting equilibrium would now involve cases in which the judiciary overturns the regulator—a possibility that was not part of the equilibrium characterized in Section 3.3.

Perhaps more significantly, however, because R's investment in analytic learning is now sunk at the time J makes its decision, from that point forward R and J share identical policy commitments, and thus J's decision to uphold or overturn R's learning plan would be noncontroversial (and in fact optimal). Thus, this type of variation on the model's assumptions effectively gives the players a collective option of abandoning the proposed reform before the experimentation period begins. While clearly desirable from social welfare grounds, this variation would do little to address R's incentive to engage in too little analytic learning ex ante.

Finally, a relatively obvious (but somewhat involved) modification of

17. Which of J's and R's preferred outcome is further from the social optimum depends on deeper parameters in the model.

the model would entail allowing J, R, and P to have policy preferences distinct from one another. Although we have suppressed this dimension in the model in order to concentrate on institutional forces, allowing actors to have differing ideologies over policy space clearly matters, and it can substantially affect equilibrium play. Given space constraints, we cannot realistically do justice to the (myriad) permutations that differing ideologies introduce. Nevertheless, to get a taste for how such issues may alter the model's predictions, suppose we altered our baseline model to allow only the judge's ideal point (y_j^*) to drift away from 0, and in the direction of the known location of the status quo ante, $\alpha > 0$. It should be clear that in the limiting case of $y_j^* = \alpha$, J will summarily reject any and all rule-making changes issued by R. Indeed, when J's ideal point coincides with the known status quo, J can do no better than to resist change of any sort. Perhaps more interesting, however, is what might happen if J's preferences occupy a middle ground between $y_j^* = 0$ and $y_j^* = \alpha$. Here, J is relatively skeptical about the desirability of reform, but not unalterably so. Anticipating J's greater reluctance, R will anticipate that the judicial review constraint $\Delta_J(\gamma^R, K^R) \geq 0$ now has become sharper, since J may not share R's views about the continuation game and, in particular, about whether the proposed experimentation period is desirable over the status quo. That knowledge can, in equilibrium, motivate R to engage in a greater degree of analytic learning than it would otherwise entertain, improving welfare. Viewed in this sense, the injection of some ideological diversity can be beneficial to social goals, even when the public does not share the views of the ideological outlier (a point we also illustrate in Spitzer and Talley [2013]). Although variations such as this are both interesting and worth pursuing, we leave them for future endeavors.

5. CONCLUSION

In this paper, we have developed a framework to study regulatory CBA within a judicial-administrative hierarchy. Our focus has been motivated by recent judicial decisions and proposed legislation in financial regulation that appear to downplay the appropriate role for regulatory experimentation relative to more conventional forms of analytic learning in conducting CBA.

Our analysis has produced three key insights. First, there do not appear to be any a priori reasons to relegate regulatory experimentation to a metaphorical backseat relative to analytic learning in CBA. Both approaches have distinct benefits and drawbacks, and thus both should be considered in administrative rule making (and the judicial review thereof). Second, we have shown that—notwithstanding the usefulness of experimentation—institutional structures can cause regulators to be too zealous about experimentalism while causing judges to be too skeptical. Thus, while we disagree with judicial decisions that (arguably) dismiss the role of experimentation, we are not surprised at the disagreement between regulators and judges about what relative weight it should receive in CBA. Finally, we have demonstrated that there is no guarantee that the threat of judicial review can promise to solve (or even substantially mollify) the institutionally grounded disagreements between regulators and the judiciary noted above. At the same time, our analysis suggests that at least some institutional structures (such as limited commitment power and some ideological diversity among actors) may provide a partial solution to the problem.

APPENDIX: PROOFS

We first present the derivation (from Section 3.1) of the distribution of hybridized signal $x = (\tau\mu + \gamma z + \omega\Sigma_{t=1}^{K} v_t)/(\tau + \gamma + \omega K)$. We show that hybrid signal $x \sim N\{\mu, (\tau\{1 + [\tau/(\gamma + \omega K)]\})^{-1}\}$. Because all components of x are normal, so must be x, and we therefore need only to compute its mean and precision. To compute the mean, we make use of the law of iterated expectations:

$$
\begin{aligned}
E(x) = E_{y_1}[E_{x|y_1}(x|y_1)] &= E_{y_1}\left[E_{xy_1}\left(\frac{\tau\mu + \gamma z + \omega\Sigma_{t=1}^{K} v_t}{\tau + \gamma + \omega K}\middle| y_1\right)\right] \\
&= \frac{\tau\mu + \gamma E_{y_1}(y_1) + \omega K E_{y_1}(y_1)}{\tau + \gamma + \omega K} = \mu.
\end{aligned}
\tag{A1}
$$

To compute the precision of x, we make use of the law of iterated variance:

$$
\begin{aligned}
\mathrm{var}(x) &= E[\mathrm{var}(x|y_1)] + \mathrm{var}[E(x|y_1)] \\
&= E\left[\mathrm{var}\left(\frac{\tau\mu + \gamma z + \omega\Sigma_{t=1}^{K} v_t}{\tau + \gamma + \omega K}\middle| y_1\right)\right] + \mathrm{var}\left(\frac{\tau\mu + \gamma y_1 + \omega\Sigma_{t=1}^{K} y_1}{\tau + \gamma + \omega K}\right) \\
&= E\left[\frac{\gamma^2 \mathrm{var}(z y_1) + \omega^2 \Sigma_{t=1}^{K} \mathrm{var}(v_t y_1)}{(\tau + \gamma + \omega K)^2}\right] + \mathrm{var}\left[\left(\frac{\gamma + \omega K}{\tau + \gamma + \omega K}\right) y_1\right] \\
&= \frac{\gamma + \omega K}{(\tau + \gamma + \omega K)^2} + \frac{(\gamma + \omega K)^2}{\tau(\tau + \gamma + \omega K)^2} = \left(\frac{1}{\tau}\frac{\gamma + K\omega}{\tau + \gamma + K\omega}\right).
\end{aligned}
\tag{A2}
$$

Taking the inverse of this variance yields precision for x of $\tau[1 + \tau/(\gamma + \omega K)]$, which is the expression from the text.

Proposition 1. If $\alpha^2 > 1/\tau$ and $c'(0)$ is sufficiently small, then $\alpha^{\mathrm{ESW}} > 0$. Similarly, if $\alpha^2 > 1/\tau$ and μ is sufficiently small, then $K^{\mathrm{ESW}} > 0$.

Proof. Recall the social planner's expected payoff over the status quo from implementing learning plan (γ, K):

$$\begin{aligned}\Delta_{\mathrm{ESW}}(\gamma, K) = {}& -c(\gamma) - (N+2)\left(\beta\frac{1-\beta^K}{1-\beta}\right)\left(\mu^2 + \frac{1}{\tau} - \alpha^2\right) \\ & - (N+2)\beta^{K+1}\frac{1-\beta^{(T-K)}}{1-\beta}E_x\left[\min\left\{0,\ x^2 + \frac{1}{\tau+\gamma+\omega K} - \alpha^2\right\}\right].\end{aligned} \tag{A3}$$

This function is strictly concave, moreover, in both γ and K. Differentiating $\Delta_{\mathrm{ESW}}(\gamma, K)$ with respect to γ and imposing the most restrictive condition of $\gamma = K = 0$ yields

$$\begin{aligned}&\left.\frac{\partial\Delta_{\mathrm{ESW}}(\gamma, K)}{\partial\gamma}\right|_{\gamma=K=0} = \\ &-c'(0) + (N+2)\beta\frac{1-\beta^T}{1-\beta}\frac{\Pr\{x \in [-\sqrt{\alpha^2 - 1/\tau},\ \sqrt{\alpha^2 - 1/\tau}]\}}{\tau^2}.\end{aligned} \tag{A4}$$

The second term is defined and is strictly positive so long as $\sqrt{\alpha^2 - 1/\tau} > 0$ (the condition stated in the proposition), and thus the entire expression is positive so long as $c'(0)$ is sufficiently small.

Now consider the second statement in the proposition. To show the second part of the proposition, hold $\gamma = 0$ and compare $\Delta_{\mathrm{ESW}}(0, 1) - \Delta_{\mathrm{ESW}}(0, 0)$:

$$\Delta_{\mathrm{ESW}}(0,\ 0) = -(N+2)(\beta)\frac{1-\beta^T}{1-\beta}\left[\min\left\{0,\ \mu^2 + \frac{1}{\tau} - \alpha^2\right\}\right],$$

$$\begin{aligned}\Delta_{\mathrm{ESW}}(0,\ 1) = {}& -(N+2)\beta\left(\mu^2 + \frac{1}{\tau} - \alpha^2\right) \\ & - (N+2)\beta^2\frac{1-\beta^{(T-1)}}{1-\beta}E_x\left[\min\left\{0,\ x^2 + \frac{1}{\tau+\omega} - \alpha^2\right\}\right].\end{aligned} \tag{A5}$$

Their difference is therefore

$$\begin{aligned}\Delta_{\text{ESW}}(0,\ 1) - \Delta_{\text{ESW}}(0,\ 0) =\ & -(N+2)\beta\Big(\mu^2 + \frac{1}{\tau} - \alpha^2\Big) \\ & -\ (N+2)\beta^2\frac{1-\beta^{(T-1)}}{1-\beta}E_x\Big[\min\Big\{0,\ x^2 + \frac{1}{\tau+\omega} - \alpha^2\Big\}\Big] \\ & +\ (N+2)(\beta)\frac{1-\beta^T}{1-\beta}\Big[\min\Big\{0,\ \mu^2 + \frac{1}{\tau} - \alpha^2\Big\}\Big].\end{aligned}$$

The condition $\alpha^2 > 1/\tau$ (stated in the proposition) implies that the first term in this expression is strictly positive as long as μ is sufficiently small (the other condition in the proposition) or equivalently that $\mu^2 + (1/\tau) - \alpha^2 < 0$. When these conditions hold, the above expression can be simplified to read

$$\Delta_{\text{ESW}}(0,\ 1) - \Delta_{\text{ESW}}(0,\ 0) =$$

$$-(N+2)\beta^2\frac{1-\beta^{(T-1)}}{1-\beta}\Big[E_x\Big(\min\Big\{0,\ x^2 + \frac{1}{\tau+\omega} - \alpha^2\Big\}\Big) - \Big(\mu^2 + \frac{1}{\tau} - \alpha^2\Big)\Big],$$

which is strictly positive so long as

$$E_x\Big[\min\Big\{0,\ x^2 + \frac{1}{\tau+\omega} - \alpha^2\Big\}\Big] - \Big(\mu^2 + \frac{1}{\tau} - \alpha^2\Big) < 0. \tag{A6}$$

It is clear that $E_x[\min\{0,\ x^2 + 1/(\tau + \omega) - \alpha^2\}] < E_x[x^2 + 1/(\tau + \omega) - \alpha^2]$, and thus by substitution into equation (A6) we can derive the following sufficient condition for equation (A6) to hold

$$E_x\Big(x^2 + \frac{1}{\tau+\omega} - \alpha^2\Big) - \Big(\mu^2 + \frac{1}{\tau} - \alpha^2\Big) \le 0 \Leftrightarrow E(x^2) + \frac{1}{\tau+\omega} - \mu^2 - \frac{1}{\tau} \le 0.$$

From the distribution derived on x above, we have

$$\begin{aligned}E(x^2) &= \operatorname{var}(x) + E(x)^2 \\ &= \Big(\frac{1}{\tau}\frac{\gamma + K\omega}{\tau + \gamma + K\omega}\Big)\Big|_{K=1;\gamma=0} + \mu^2 = \frac{1}{\tau}\frac{\omega}{\tau+\omega} + \mu^2.\end{aligned}$$

And thus we have

$$E(x^2) + \frac{1}{\tau+\omega} - \mu^2 - \frac{1}{\tau} = \frac{1}{\tau}\frac{\omega}{\tau+\omega} + \mu^2 + \frac{1}{\tau+\omega} - \mu^2 - \frac{1}{\tau} = 0,$$

which is clearly nonpositive, and it therefore follows that equation (A6) is satisfied as a strict inequality under the conditions stated in the proposition. QED.

Proposition 2. The regulator prefers to engage in (weakly) too little analytic learning and (weakly) excessive field experimentation relative to the social optimum. However, the regulator would implement a socially optimal learning protocol if it were constrained to choose any $\gamma \geq \gamma^{ESW}$.

Proof. The proposition is a direct implication of differentiating to equation (9) to equation (8) with respect to γ. QED.

Proposition 3. The judiciary prefers to engage in (weakly) too much analytic learning and (weakly) too little field experimentation relative to the social optimum.

Proof. The proposition is a direct implication of differentiating equation (10) to equation (8) with respect to K. QED.

Proposition 4. There is a unique Bayesian perfect equilibrium in the sequential game defined above. In it, the regulator announces a learning plan (γ^R, K^R) that maximizes its own welfare $\Delta_R(\gamma, K)$ unconstrained by the specter of judicial review. As per proposition 2, this plan entails (weakly) too little analytic learning and (weakly) too much field experimentation relative to the social optimum. This policy is never overturned by the judiciary.

Proof. The proof is by construction. Suppose that R solved programming problem (11) but ignored the third constraint in that program (which requires $\Delta_J(\gamma, K) \geq 0$). Let the solution to this problem be denoted (γ^R, K^R). Because (γ^R, K^R) is optimal for the regulator by hypothesis, then it must be the case that $\Delta_R(\gamma^R, K^R) \geq 0$. However, comparing equation (9) with equation (10), it is clear that because only player R bears the cost of analytic learning, $\Delta_J(\gamma^R, K^R) \geq \Delta_R(\gamma^R, K^R)$ for all (γ^R, K^R). QED.

REFERENCES

Bartlett, Robert P., III. 2014. The Institutional Framework for Cost-Benefit Analysis in Financial Regulation: A Tale of Four Paradigms? *Journal of Legal Studies* 43:S379–S405.

Berry, Donald, and Bert Fristedt, eds. 1985. *Bandit Problems: Sequential Allocation of Experiments*. Dordrecht: Springer.

Cox, James D., and Benjamin J. C. Baucom. 2012. The Emperor Has No Clothes: Confronting the DC Circuit's Usurpation of Rulemaking Authority. *Texas Law Review* 90:1811–47.

Forstall, Tim. 2012. Posner's FDA for Financial Innovation: A Truly Terrible Idea. *Forbes*. February 19. http://www.forbes.com/sites/timworstall/2012/02/19/posners-fda-for-financial-innovation-a-truly-terrible-idea/.

Gersen, Jacob. 2007. Temporary Legislation. *University of Chicago Law Review* 74:247–98.

Gubler, Zachary. 2014. Experimental Rules. *Boston College Law Review* 55: 129–79.

Kraus, Bruce, and Connor Raso. 2013. Rational Boundaries for SEC Cost-Benefit Analysis. *Yale Journal on Regulation* 30:289–342.

Lee, Alex. 2013. An Options Approach to Agency Rulemaking. *Administrative Law Review* 65:881–943.

McGinnis, John O. 2013. *Accelerating Democracy: Transforming Governance through Technology.* Princeton, N.J.: Princeton University Press.

Posner, Eric, and Glen Weil. 2013. An FDA for Financial Innovation: Applying the Insurable Interest Doctrine to Twenty-First-Century Financial Markets. *Northwestern University Law Review* 107:1307–58.

Ribstein, Larry. 2010. The SEC vs. Shareholders. *Truth on the Market* (blog), August 30. http://truthonthemarket.com/2010/08/30/the-sec-vs-shareholders/.

Romano, Roberta. 2012. Regulating in the Dark. Pp. 103–6 in *Regulatory Breakdown*, edited by Gary Coglianese. Philadelphia: University of Pennsylvania Press.

Sabel, Charles F., and William H. Simon. 2011. Minimalism and Experimentalism in the Administrative State. *Georgetown Law Journal* 100:53–93.

Spitzer, Matthew L., and Eric Talley. 2000. Judicial Auditing. *Journal of Legal Studies* 29:649–83.

———. 2013. Left, Right, and Center: Strategic Information Acquisition and Diversity in Judicial Panels. *Journal of Law, Economics, and Organization* 29: 638–80.

Sunstein, Cass. 2002. *The Cost-Benefit State: The Future of Regulatory Protection.* Chicago: American Bar Association.

———. 2014. Financial Regulation and Cost-Benefit Analysis: A Comment. Harvard Public Law Working Paper No. 14-17. Harvard University Law School, Cambridge, Mass.

Talley, Eric 2010. Perspective on Proxy Access. *The Conglomerate: Business, Law, Economics, Society* (blog). August 26. http://www.theconglomerate.org/2010/08/viewed-through-the-lens-of-editorial-pages-wednesdays-rule-changewas-a-watershed-event-shareholder-activists-have-been-c.html.

Whitehead, Charles. 2012. The Goldilocks Approach: Financial Risk and Staged Regulation. *Cornell Law Review* 97:1267–1308.

Yoon, Albert. 2003. Love's Labor's Lost: Judicial Tenure among Lower Federal Court Judges, 1945–2000. *California Law Review* 90:1029–60.